arXiv:1411.7481v1 [stat.ME] 27 Nov 2014

Nonparametric Bayesian Inference for Mean Residual Life Functions in Survival Analysis

VALERIE POYNOR[1] and ATHANASIOS KOTTAS[1]

[1] *University of California, Santa Cruz*

December 1, 2014

SUMMARY

Modeling and inference for survival analysis problems typically revolves around different functions related to the survival distribution. Here, we focus on the mean residual life function which provides the expected remaining lifetime given that a subject has survived (i.e., is event-free) up to a particular time. This function is of direct interest in reliability, medical, and actuarial fields. In addition to its practical interpretation, the mean residual life function characterizes the survival distribution. We review key properties of the mean residual life function and investigate its form for some common distributions. We then develop general Bayesian nonparametric inference for mean residual life functions built from a Dirichlet process mixture model for the associated survival distribution. Particular emphasis is placed on the selection of the mixture kernel to ensure desirable properties for the mean residual life function arising from the mixture distribution. We advocate for a mixture model with a gamma kernel and dependent baseline distribution for the Dirichlet process prior. The empirical performance of our modeling technique is studied with two simulation examples, a data set of two experimental groups, and a data set involving right censoring. Moreover, to illustrate the practical utility of the nonparametric mixture model, we compare it in the context of one of the data examples with an exponentiated Weibull model, a parametric survival distribution that allows various shapes for the mean residual life function.

KEYWORDS: Bayesian nonparametrics; Dirichlet process mixture models; Mean residual life; Right censoring; Survival function.

1. INTRODUCTION

Survival data describe the time to a particular event. This event is usually referred to as the failure of some machine or death of a person. However, survival data can also represent duration of unemployment, life expectancy of a product, the time until a patient relapses, etc. The survival function of a positive random variable T defines the probability of survival beyond time t, $S(t) = Pr(T > t) = 1 - F(t)$, where $F(t)$ is the distribution function. The hazard rate function computes the probability of a failure in the next instant given survival up to time t,

$$h(t) \quad = \quad \lim_{\Delta t \to 0} \frac{Pr[t < T \leq t + \Delta t | T > t]}{\Delta t} = \frac{f(t)}{S(t)}$$

with the expression in terms of the density function, $f(t)$, holding for continuous T. The mean residual life (mrl) function at a specific time point t computes the expected remaining survival time of a subject given survival up to time t. Assuming that $F(0) = 0$ and $\mu \equiv E(T) = \int_0^\infty S(t)dt < \infty$, the mrl function for continuous T is defined as:

$$m(t) \quad = \quad E(T - t | T > t) = \frac{\int_t^\infty (u - t) f(u) du}{S(t)} = \frac{\int_t^\infty S(u) du}{S(t)} \tag{1}$$

and $m(t) \equiv 0$ whenever $S(t) = 0$. The mrl function is of particular interest because of its easy interpretability and large area of application (Guess & Proschan, 1985). Moreover, it characterizes the survival distribution via the Inversion Formula (2). Again for continuous T with finite mean, the survival function is defined through the mrl function:

$$S(t) = \frac{m(0)}{m(t)} exp\left[- \int_0^t \frac{1}{m(u)} du \right]. \tag{2}$$

One point of interest in survival and reliability analysis is the study of the form for the mrl function of various distributions. In Section 2, we provide a discussion on some key probabilistic properties and defining characteristics of the mrl function as well as investigate the form of the mrl function under a number of common distributions. We show that the shape of the mrl function is often limited to monotonically increasing (INC) or decreasing (DCR) functions, which may be appropriate for some situations, but not suitable for other populations. For instance, biological lifetime data tend to support lower mrl during infancy and elderly age while there is a higher mrl during the middle ages. The shape of this mrl function is unimodal and commonly referred to as upside-down bathtub (UBT) shape. There have been several papers that have investigated the form of the mrl function in relation to the hazard function. A well-known relationship for monotonically increasing (decreasing) hazard functions is that the corresponding mrl function will

be monotonically decreasing (increasing); see Finkelstein (2002) for a review. Gupta & Akman (1995) establish sufficient conditions for the mrl function to be decreasing (increasing) or UBT (BT) given that the hazard is BT (UBT). Xie et al. (2004) look at the specific change points of mrl function and hazard function.

Another point of interest lies in inference for the mrl function. The classical survival analysis literature includes several estimation techniques for mrl functions. The most basic estimator, being the empirical estimate, was first studied in Yang (1978). The empirical estimate is defined by $\hat{m}_n(t) = ((\int_t^\infty S_n(u)du)/S_n(t))\delta_{[0,T_{(n)}]}(t)$ where $S_n(t)$ is the empirical survival function and $T_{(n)}$ is the maximum observed survival time. It is shown that under this fixed finite interval, the estimator is asymptotically unbiased, is uniformly strong consistent, and as n goes to infinity it converges in distribution to a Gaussian process. Hall & Wellner (1979) extend the empirical estimator by defining it for values on the positive real line. Furthermore, they provide nonparametric confidence bands for the estimate via transformations of the limiting process of the estimator into Brownian motion. Abdous & Berred (2005) use a local linear fitting technique to find a smooth estimate assuming only that the smoothing kernel is symmetric. A nonparametric hypothesis testing procedure for comparing mrl functions from two independent groups was introduced by Berger et al. (1988). A practical benefit of this procedure is that mrl estimates of the two groups were allowed to cross, a pattern that is likely to arise in applications.

Classical estimation for the mrl function began to have a semiparametric regression flavor when Oakes & Dasu (1990) extended the class of distribution having linear mrl functions (Hall & Wellner, 1981), to a family having proportional mrl functions, $m_1(t) = \psi m_2(t)$ for $\psi > 0$. Maguluri & Zhang (1994) further extended the proportional mrl model to a regression setting, $m(t|z) = exp(\psi z)m_0(t)$, where z is a vector of covariates, ψ is of vector of regression coefficients, and $m_0(t)$ is a baseline mrl function. Chen & Cheng (2005) also extend the proportional mrl model to include inference for the regression parameters with censored data.

In contrast to the classical literature, there has been very little work on modeling and inference for mrl functions under the Bayesian framework. Lahiri & Park (1991) present nonparametric Bayes and empirical Bayes estimators under a Dirichlet process (Ferguson, 1973) prior for the probability distribution. They show that the Bayes estimator becomes a weighted average of the prior guess for the mrl function and the empirical mrl function of the data. Johnson (1999) discusses a Bayesian method for estimation of the mrl function under interval and right censored data, also using a Dirichlet process prior for the corresponding survival function.

In Section 3, we develop Bayesian inference methods for mrl functions using a Dirichlet process mixture modeling approach. A detailed discussion on the importance of kernel selection is

conveyed. In particular, we provide a sufficient condition that ensures existence and finiteness of the mrl function for any given mixture kernel and Dirichlet process prior baseline distribution. In addition, we study the effect of the kernel choice on the tail behavior of the mrl function of the mixture distribution. In Section 4, we provide a number of illustrations of our modeling techniques, including comparison of two experimental groups and modeling in the presence of right censoring. Concluding remarks and possible extensions to more general modeling for mrl functions are discussed in Section 5.

2. THEORY AND PROPERTIES OF MRL FUNCTIONS

In this section we review some important properties and characteristics of the mrl function. We move on to discuss the shapes of mrl functions, and provide the form of the mrl function for several common distributions.

2.1 *Properties of mrl functions*

We start out by recalling an elementary relationship between the survival function and the moments of the distribution. If the r^{th} moment exists for a continuous random variable T, we have:

$$E(T^r) \quad = \quad r \int_0^\infty t^{r-1} S(t) dt \tag{3}$$

This expression is of interest, because once the Inversion Formula (2) is applied, we have a way of obtaining the moments (when they exist) from the mrl function.

The derivation of the expression for the mrl function through the survival function as stated in (1), can be found in Smith (2002). Using (1), we can easily confirm that the first moment is equivalent to the mrl function at $t = 0$.

$$m(0) = \frac{\int_0^\infty S(u)du}{S(0)} = \int_0^\infty S(u)du = \mu \tag{4}$$

Similar to the moment generating function of probability distributions, when the mrl function exists, it defines the survival distribution. We can obtain the survival function using the Inversion Formula (2). The derivation of (2) can be found in Smith (2002). Hall & Wellner (1981) provide important properties of the mrl function for the development of the characterization theorem for mrl functions. The characterization theorem states the necessary and sufficient conditions such that a function is the mrl function for a survival distribution. For our purposes, the defining

characteristics of interest are that $m(t)+t$ is non-decreasing and survival distribution has finite mean, $\mu = m(0) < \infty$.

2.2 *Forms of mrl functions*

The most basic shape of the mrl function is a line. Oakes & Dasu (1990) focus on linear mrl functions discussed in Hall & Wellner (1981). The key result is that if the mrl function is linear, $m(t) = At+B$ $(A > -1, B > 0)$, then by use of the Inversion Formula (2), the survival function has the form, $S(t) = [B/(At+B)]_+^{1/A+1}$. When $A = 0$, the survival distribution is an exponential with mean β. For $A > 0$ the survival function is a Pareto distribution under the linear transformation, $Z = AT + B$, with shape parameter $\alpha = 1/A + 1$ and scale parameter $\beta = B$. For $-1 < A < 0$ the survival function is a rescaled beta distribution under the transformation, $Z = -AT$. The pdf of a rescaled beta distribution is $f(z; a, b, p, q) = ((z-a)^{p-1}(b-z)^{q-1})/(\mathbf{B}(p,q)(b-a)^{p+q+1})$ where $a \leq z \leq b,\ p, q > 0$, and $\mathbf{B}(.,.)$ is the beta function, specifically, $B = b, a = 0, p = 1, q = -(1/A+1)$.

These linear mrl functions are convenient in the sense that they yield a closed form for the mrl function, however, the linearity is too limiting to be of much practical use. There are a number of distributions having more flexible mrl functions, such as increasing and decreasing curvatures as well as BT or UBT shapes. The difficulty for these distributions lies in obtaining a closed form of the mrl function.

The distributions discussed here have no known closed form for their associated mrl making them difficult to explore. However, through the use of (5), see Govilt & Aggarwal (1983) for derivation, and/or simple transformations of T, we are able to obtain forms of the mrl functions comprised of well-known integrals, that are easily evaluated with most statistical programming software.

$$m(t) \quad = \quad \frac{\int_t^\infty u f(u) du}{S(t)} - t \tag{5}$$

Govilt & Aggarwal (1983) use (5) to obtain the mrl form, $m(t) = \left(t^\alpha exp\left[-t/\lambda\right]\right) / \left(\lambda^{\alpha-1}\Gamma(\alpha) S_{\mathbf{T}}(t)\right) + \lambda\alpha - t$, for the gamma distribution with shape parameter α and scale parameter λ. The mrl function of the Gompertz distribution with shape and scale parameters $\alpha, \lambda > 0$ respectively can be simplified using (1) and the transformation $z(t) = (\lambda/\alpha)e^{\alpha t}$ to give$m(t) = e^{z(t)}\,(1/\alpha)\,\Gamma_{inc}(0, z(t))$. The loglogistic distribution with shape and scale parameters $\alpha > 1, \lambda > 0$ can be written as $m(t) = (\lambda/\alpha)\,\Gamma\left(1-(1/\alpha)\right)\Gamma\left(1/\alpha\right) S_{\mathbf{Z}}\left(z(t); 1-(1/\alpha), (1\alpha)\right)\left(1+(t/\lambda)^\alpha\right)$ by simplifying (1) as is done by Gupta et al. (1999). The mrl of the lognormal distribution, with location μ and scale σ^2, can be simplified to be $m(t) = \left(e^{\left(\mu+(\sigma^2/2)\right)}\left[1-\Phi\left((ln(t)-(\mu+\sigma^2))/\sigma\right)\right]\right) / \left(1-\Phi\left(ln(t)-\mu\right)/\sigma\right) - t$ through

(5). Using (1) and the transformation $z(t) = t^{\alpha}$, the mrl function of the Weibull distribution with shape parameter $\alpha > 0$ and scale parameter $\lambda > 0$ has form $m(t) = \left((\lambda/\alpha)\,\Gamma\left(1/\alpha\right) S_{\mathbf{Z}}(z(t))\right)/S_{\mathbf{T}}(t)$.

Table 1, provides a summary of the shapes of the hazard rate and mrl functions for these distributions. The table shows how restricted these commonly used distribution are in modeling the mrl function. The gamma and Weibull are more versatile as they offer three potential shapes for the mrl function, but none of these shapes consider change points in the mrl function.

Distribution	Density Function	Mean Residual Life
Gamma(γ, β) shape $\gamma > 0$, rate $\beta > 0$	$\frac{\beta^{\gamma}}{\Gamma(\gamma)} t^{\gamma-1} e^{-\beta t}$	$\gamma < 1$ INC, $\gamma = 1$ constant (β^{-1}), $\gamma > 1$ DCR
Gompertz(γ, λ) shape $\gamma > 0$, scale $\lambda > 0$	$\lambda\gamma e^{\lambda t} e^{\gamma} exp\left(-\gamma e^{\lambda t}\right)$	$\forall\gamma$ DCR
Loglogistic(γ, λ) shape $\gamma > 0$, scale $\lambda > 0$	$\frac{(\gamma/\lambda)(t/\lambda)^{\gamma-1}}{[1+(t/\lambda)^{\gamma}]^2}$	$\gamma \leq 1$ undefined, $\gamma > 1$ BT
Lognormal(μ, σ) location $\mu \in \mathbb{R}$, scale $\sigma^2 > 0$	$\frac{1}{t\sqrt{2\pi\sigma^2}} e^{-\frac{(ln(t)-\mu)^2}{2\sigma^2}}$	BT
Weibull(γ, λ) shape $\gamma > 0$, scale $\lambda > 0$	$\frac{\gamma}{\lambda}\left(\frac{t}{\lambda}\right)^{\gamma-1} e^{-(t/\lambda)^{\gamma}}$	$\gamma < 1$ INC, $\gamma = 1$ constant (λ), $\gamma > 1$ DCR

Table 1: Shapes of the mean residual life function for common parametric distributions. Shapes are described as being increasing (INC), decreasing (DCR), upside down bathtub (UBT), bathtub (BT), constant, or undefined.

Modifications of the Weibull model have been explored in order to develop a more flexible parametric model with regard to the shapes of the hazard and mrl functions; see Pham & Lai (2007) for an extensive list. We focus of the exponentiated Weibull distribution which has closed form survival function and can take on a number of various forms for the mrl (Mudholkar & Strivasta, 1993); see Table 2. The distribution and mrl functions for the exponentiated Weibull model are given by the following expressions:

$$\begin{aligned} F(t;\alpha,\theta,\sigma) &= \left[1 - exp\left(-\left(\frac{t}{\sigma}\right)^{\alpha}\right)\right]^{\theta}, \quad t > 0, \alpha, \theta, \sigma > 0 \\ m(t;\alpha,\theta,\sigma) &= \frac{\int_t^{\infty}\left[1 - \left[1 - exp\left(-\left(\frac{u}{\sigma}\right)^{\alpha}\right)\right]^{\theta}\right] du}{1 - \left[1 - exp\left(-\left(\frac{t}{\sigma}\right)^{\alpha}\right)\right]^{\theta}} \end{aligned} \tag{6}$$

where α and θ are shape parameters and σ is a scale parameter. Note that σ, being a scale, will not play a role in determining the form of the hazard and mrl functions. Table 2 provides the parameter sets that result in each distinct shape for the mrl function.

Mudholkar & Strivasta (1993) provide a similar table as Table 2 for the hazard rate function for specific domains of α and θ. Xie et al. (2004) look at the role of the product of the shape

α	θ	$\alpha\theta$	form of mrl function
1	1	1	exponential distribution $\rightarrow$ constant mrl
–	1	–	weibull distribution $\rightarrow$ monotone (inc, dcr or constant) mrl
< 1	$\neq 1$	< 1	increasing
> 1	$\neq 1$	> 1	decreasing
> 1	< 1	< 1	UBT
< 1	> 1	> 1	BT

Table 2: Forms of MRL for Exponentiated Weibull Distribution

parameters on the form of the hazard rate. Gupta & Akman (1995) prove that if the hazard rate function is BT and $h(0) > 1/m(0)$, then the mrl is UBT, while $h(0) \leq 1/m(0)$ implies decreasing mrl function. Similarly if the hazard rate function is UBT and $h(0) > 1/m(0)$, then the mrl is BT, while $h(0) \geq 1/m(0)$ implies increasing function. Combining the aforementioned results, we can improve the table in Mudholkar & Strivasta (1993) to specify the exact shape of the mrl function for particular values of α and θ in conjunction with the value of the product of the parameters.

3 NONPARAMETRIC MIXTURE MODEL FOR MRL INFERENCE

In this section, we discuss our modeling methods for obtaining inference for the mrl function. Section 3.1 motivates the use of a nonparametric Dirichlet process mixture model (DPMM). We provide the model structure, and discuss the choice of kernel distribution. In Section 3.2, we discuss prior specification. Section 3.3 provides the techniques used to obtain posterior inference for the mixture distribution and functionals thereof.

3.1 *Model formulation*

When the data exhibits unusual distributional features such as multi-modality or skewness, parametric models tend to fail to capture these important features. A way to go about this issue is to use a mixture model that combines a number of distributions that we will refer to as components of the model. The question then becomes how many components should be used and how should they be combined together? These concerns can be addressed by bringing in a nonparametric aspect to the model, in particular, to the weights of each component and to the number of components.

We use a Dirichlet process (DP) prior for the mixing distribution resulting in a DP nonparametric mixture model, $f(t; G) = \int k(t; \boldsymbol{\theta}) dG(\boldsymbol{\theta})$, for the density of the survival distribution. In practice, an appropriately supported kernel distribution, $k(t; \boldsymbol{\theta})$, is selected, and a $DP(\alpha, G_0)$ prior

is assigned to G. The DP is a stochastic process with random sample paths that are distributions (Ferguson, 1973). Thus a realization from the DP provides a random cdf sample path. The G_0 parameter is the baseline or centering distribution, while α is a precision parameter; the larger the value of α the closer the DP sample path is to the centering distribution. We use the stick-breaking (SB) constructive definition of the DP defined by Sethuraman Sethuraman (1994), which states that a sample $G(\cdot)$ from $DP(\alpha, G_0)$ is almost surely of the form $\sum_{l=1}^{\infty} w_l \delta_{\boldsymbol{\theta}_l}(\cdot)$ where $\delta_{\boldsymbol{\theta}_l}(\cdot)$ is a point mass at $\boldsymbol{\theta}_l$. The $\boldsymbol{\theta}_l$, for all $l \in \{1, 2, ...\}$, are iid samples from the baseline distribution, G_0, and the w_l are the corresponding weights constructed sampling iid latent variables $v_r \sim$Beta$(1, \alpha)$, for all $r \in \{1, 2, ...\}$, then $w_1 = v_1$ and $w_l = v_l \prod_{r=1}^{l-1}(1 - v_r)$, for $l \in \{2, 3, ...\}$.

We use the truncated version of the SB constructive definition of the DP, $G_L(\cdot) = \sum_{l=1}^{L} p_l \delta_{\boldsymbol{\theta}_l}(\cdot)$, where $\boldsymbol{\theta}_l \overset{iid}{\sim} G_0$ for $l = 1, ..., L$, and $p_1 = v_1$, $p_l = v_l \prod_{r=1}^{l-1}(1 - v_r)$, where $v_r \overset{iid}{\sim} Beta(1, \alpha)$ for $r = 1, ..., L - 1$, and $p_L = 1 - \sum_{l=1}^{L-1} p_l$. The model is given by:

$$f(t; G) \sim \int k(t; \boldsymbol{\theta}) dG(\boldsymbol{\theta}) = \sum_{l=1}^{L} p_l k(t; \boldsymbol{\theta}_l) \tag{7}$$

where p_l for $l = 1, ..., L$ are the weights obtained via the SB construction, described above, corresponding to the component $\boldsymbol{\theta}_l$ and L is the total number of components in the mixture model. Technically, since the number of components is predetermined there is no nonparametric element to the number of components. However, L is generally chosen to overestimate the true number of components, so that the number of components suggested by the data is captured by the model. In fact, many of the components will just be assigned a probability that is virtually zero. The number of components for the finite sum DP approximation can be found using $E(\sum_{l=1}^{L} p_l) = 1 - (\alpha/(\alpha+1))^L$, in particular, solving for L in $(\alpha/(\alpha + 1))^L = \epsilon$ for small $\epsilon > 0$.

Our primary aim in this paper is to present a Bayesian model that provides both flexible and practical inference for the mean residual life function. The mrl function is defined by the distribution function and vice versa, thus we advocate for the nonparametric Dirichlet process mixture which provides flexible modeling on the distribution function. We obtain inference for the mrl function via fitting a DPMM on the distribution function. Since our interest is inference for the mrl function, it is necessary that the mrl function of the DPMM exists and is finite. A sufficient condition for the finiteness of the mrl function for a given kernel distribution is provided later in this section. Although we do not place a prior directly on the mrl function, from the lemma stated at the end of this section, we can use prior knowledge of the tail behavior to select an agreeable kernel distribution. Essentially, we can induce a prior for the mrl function through the tail behavior. We complete the model formulation by addressing the aspect of dependency within $\boldsymbol{\theta}$. We consider

modeling the dependence between the kernel parameters by using a joint G_0 distribution in the DPMM.

Care is needed in selecting a kernel distribution to ensure the mean of the DPMM is finite, $E(T;G) < \infty$ where $T \sim f(t;G) = \int_\Theta k(t;\boldsymbol{\theta})dG(\boldsymbol{\theta})$. We provide sufficient conditions to ensure finiteness of the mean by following the argument in Theorem 3 of Ferguson (1973). Let $Z = E(T;G)$, where T is a $\mathbb{R}^+$ random valued. Recall that if $E(Z) < \infty$, then $Z < \infty$ almost surely. Hence we need to show $E(Z) < \infty$. Observe that $Z = \int_T t \int_\Theta k(t;\boldsymbol{\theta})dG(\boldsymbol{\theta})dt = \int_\Theta E(T;\boldsymbol{\theta})dG(\boldsymbol{\theta}) = \sum_{j=1}^\infty w_j W(\boldsymbol{\theta}_j)$ where the w_j are the weights arising from the stick-breaking process and $W(\boldsymbol{\theta}_j) = E(T;\boldsymbol{\theta}_j)$. Define the sequence of $\mathbb{R}^+$ valued random variables $Z_n = \sum_{j=1}^n w_j W(\boldsymbol{\theta}_j)$, for $n = 1, 2, ...$ Note that Z_n is an almost surely increasing sequence and $Z_n \overset{a.s.}{\to} Z$. Thus by the monotone convergence theorem, $E(Z_n) \overset{a.s.}{\to} E(Z)$. Now, we can write $E(Z) = E(\sum_{j=1}^\infty w_j W(\boldsymbol{\theta}_j))$. Using the independence of w_j and $W(\boldsymbol{\theta}_j)$, the expectation becomes $\sum_{j=1}^\infty E(w_j)E(W(\boldsymbol{\theta}_j)) = \sum_{j=1}^\infty E(w_j)(\int_{\boldsymbol{\Theta}} W(\boldsymbol{\theta}_j)dG_0(\boldsymbol{\theta}_j))$. Upon integration over $\boldsymbol{\theta}_j$ in the last expression, the resulting expression is free of the subscript j and is a function of the parameters of the baseline distribution, G_0, with parameters $\boldsymbol{\psi}$. Define $A(\boldsymbol{\psi}) = \int_{\boldsymbol{\Theta}} W(\boldsymbol{\theta}_j)dG_0(\boldsymbol{\theta}_j)$, then $E(Z)$ becomes $A(\boldsymbol{\psi}) \sum_{j=1}^\infty E(w_j) = A(\boldsymbol{\psi})$, since $\sum_{j=1}^\infty E(w_j) = 1$. Therefore, if $A(\boldsymbol{\psi}) = \int_{\boldsymbol{\Theta}} W(\boldsymbol{\theta})dG_0(\boldsymbol{\theta}) < \infty$, then $E(T;G) < \infty$ almost surely. In words, the finiteness of the expected value of the mean of the kernel distribution with respect to the baseline distribution guarantees finiteness of the first moment of the DPMM.

Common kernel distributions in mixture modeling for survival data include the lognormal, Weibull, and gamma distributions (e.g., Kuo & Mallick, 1997; Kottas, 2006; Hanson, 2006). First, consider the lognormal kernel with $W(\boldsymbol{\theta}) \equiv E(T;\mu,\sigma) = exp(\mu+\sigma^2/2)$ and $G_0 = N(\mu;\lambda,\tau^2)\Gamma^{-1}(\sigma^2; a, \rho$ (scale)). Here, we have $A(\boldsymbol{\psi}) = \int_{-\infty}^\infty exp(\mu)N(\mu;\lambda,\tau^2)d\mu \int_0^\infty exp(\sigma^2/2)\Gamma^{-1}(\sigma^2;a,\rho)d\sigma^2$. The first integral is clearly finite, but the second integral would require a bound on σ^2 that would depend on ρ in order to be finite. We can get around the restriction by using a gamma baseline distribution, but the rate parameter of the gamma distribution would have to be truncated below at 1/2. In either case, we will not have conjugacy.

If we use a Weibull kernel with $W(\boldsymbol{\theta}) \equiv E(T;\gamma,\sigma) = \sigma^{1/\gamma}\Gamma(1+1/\gamma)$ and $G_0 = \Gamma(\gamma;a,\rho$ (rate))$\Gamma^{-1}(\sigma;c,\lambda)$, $A(\boldsymbol{\psi})$ is given by $\int_0^\infty \int_0^\infty \sigma^{1/\gamma}\Gamma(\gamma;a,\rho)\Gamma^{-1}(\sigma;c,\lambda)d\sigma d\gamma$. We can integrate out σ without difficulty by recognizing another gamma distribution, however, the finiteness of the first integral requires $\gamma > 1/c$. This is not an unreasonable restriction for $c > 1$, allowing for decreasing and/or unimodal components in the mixture, however, the second integral yields more restriction. We can obtain finiteness by constructing a function, $g(\gamma)$ that is greater than $\Gamma(1+1/\gamma)\Gamma(c-1/\gamma)\lambda^{1/\gamma}$ for $\gamma > 1/c$, where the second part of the expression is a result of the first integral. Note that $\Gamma(c-1/\gamma)$ goes to ∞ as $\gamma \to 1/c$, which causes problems in conver-

gence. We can get around this by making γ bounded below by a value just slightly larger than $1/c$. By showing $E(g(\gamma)) < \infty$ with respect to the distribution $\Gamma^{-1}(\gamma; a, \rho)$, then we know that $E(\Gamma(1+1/\gamma)\Gamma(c-1/\gamma)\lambda^{1/\gamma}) < \infty$. The idea is that $E(g(\gamma))$ is easily computed, so it is convenient to use a function of the form $g(\gamma) = \lambda^v(1/\gamma^w + \Gamma(c))$ for $v, w > 0$. Using this function form, will result in a restriction on the shape parameter, a, that will depend on w. Since w and a are both fixed parameters, this is not an unreasonable restriction, and slice sampling may be used in the MCMC for γ, but the sampling from the posterior conditional of ρ will require a Metropolis-Hastings step.

Consider a gamma kernel distribution with $W(\boldsymbol{\theta}) \equiv E(T; \alpha, \beta) = \alpha/\beta$ and $G_0 = f(\alpha; \boldsymbol{\omega})\Gamma(\beta; c, \lambda)$. We can separate the integrals in $A(\boldsymbol{\psi})$ to be $\int_0^\infty \alpha f(\alpha; \boldsymbol{\omega}) d\alpha \int_0^\infty \beta^{-1}\Gamma^{-1}\beta^{-1}\Gamma(\beta; c, \lambda) d\beta$, where the first integral is simply $E(\alpha)$ with respect to $f(\alpha; \boldsymbol{\omega})$ and the second integral is $E(\beta^{-1})$ where $\beta^{-1} \sim \Gamma^{-1}(\beta; c, \lambda)$. Therefore as long as we choose $f(\alpha; \boldsymbol{\omega})$ to have finite mean and set $c > 1$, then $A(\boldsymbol{\psi}) < \infty$. We do not have conjugacy in the MCMC for α, but our parameter restriction is minimal. The mean and variance of the gamma distribution are not independent, so we might consider a joint G_0. A convenient option would be to model using $\boldsymbol{\theta} = (\theta = log(\alpha), \phi = log(\beta))'$ and place a bivariate normal distribution on $G_0 = N_2(\boldsymbol{\theta}; \boldsymbol{\mu}, \boldsymbol{\Sigma})$. Now, we have $A(\Psi) = \int\int e^{\theta-\phi} N_2((\theta, \phi)'; \boldsymbol{\mu}, \boldsymbol{\Sigma}) d\theta d\phi = E(e^{\theta-\phi})$ which is the moment generating function of the bivariate normal, $E(e^{\boldsymbol{\theta}' t})$, with $t = (1, -1)'$. Hence $A(\Psi) = e^{(1,-1)\boldsymbol{\mu} + (1/2)(1,-1)\boldsymbol{\Sigma}(1,-1)'}$, which is finite for any $\boldsymbol{\mu} \in \mathbb{R}^2$ and any non-negative definite $\boldsymbol{\Sigma} \in \mathbb{R}^{2x2}$.

Another important consideration in the choice of the mixture kernel is the shape of the mixture mrl function relative to the mrl function of the kernel distribution. The following lemma, whose proof can be found in Appendix A.1, provides a result on the tail behavior of the mrl function for the mixture distribution.

Lemma 1. Let $m(t; \boldsymbol{\theta})$ be the parametric mrl function of the corresponding to the DPMM kernel and $m(t; G_L)$ be the mrl function of the mixture, where G_L is the truncated approximation to the mixing distribution. Then,
1. If $\lim_{t\to\infty} m(t; \boldsymbol{\theta}) = \infty \; \forall \boldsymbol{\theta} \in \boldsymbol{\Theta}$, then $\lim_{t\to\infty} m(t; G_L) = \infty$.
2. If $\lim_{t\to\infty} m(t; \boldsymbol{\theta}) = 0 \; \forall \boldsymbol{\theta} \in \boldsymbol{\Theta}$, then $\lim_{t\to\infty} m(t; G_L) = 0$.

Taking into account the condition for $E(T; G) < \infty$ and the lemma above, the gamma distribution emerges as the more suitable choice for the kernel distribution. Referring back to Table 1, we can see that a lognormal kernel will always result in a mrl that goes to infinity in tail. A Gompertz kernel would result in a mrl that tends to zero in the tail. If there is prior knowledge regarding the tail behavior of the mrl, then it would make sense to choose a kernel that has a

corresponding mrl with agreeable tail behavior. However, in the case that prior knowledge of the mrl tail behavior is not known, the gamma or Weibull kernel would be appropriate choices. Per our discussion regarding the sufficient condition for existent and finite mrl, the Weibull requires more restriction on the support of the model parameters.

Another important model property to investigate is that of denseness. Let $\mathscr{F}$ represent the space of absolutely continuous distribution functions on $\mathbb{R}^+$ with finite mean. Formally, a class of distributions, $\mathcal{C}$, is said to be dense in $\mathscr{F}$, if for any distribution function, F, there exists a sequence of distribution functions, $\{F_n\} \subseteq \mathcal{C}$, that converges to F. The type of convergence implies a measure of distance between the limiting sequence and F. Johnson & Taaffe (1998) show denseness of infinite and finite mixtures of Erlang distributions. They provide details for weak convergence and make an argument for uniform convergence. A mixture of Erlang distributions is in the class of gamma mixtures, so the result holds true for gamma mixtures as well. More interesting from our prospective, however, is the denseness of the resulting mrl function. In Appendix A.2, we show that for any continuous mrl function, $m(t)$, there exists a corresponding sequence of mrl functions for a mixture of gamma distributions, $\{m_n(t)\}$ such that for any $t_0 \geq 0$, $\lim_{n\to\infty} m_n(t_0) = m(t_0)$, converges pointwise, providing the denseness result in Lemma 2.

Lemma 2. The set of mrl functions corresponding with the class of gamma mixture distributions is dense, in the pointwise sense, in the space of continuous mrl functions.

Finally we turn to the choice of G_0. We seek to be more general in our modeling by using a dependent G_0 for the parameters of the gamma kernel. This allows the model to capture correlations between the kernel parameters. Note that, once one leaves the setting of normal mixtures, the kernel parameters are not naturally separated as location and scale parameters, making the assumption of an independent G_0 more restrictive than in mixing with Gaussian kernels. Recall that modeling the shape and rate parameters of the gamma kernel on the log-scale allow us to use a bivariate normal for G_0, and we only need a non-negative definite 2×2 covariance matrix in G_0 to satisfy the sufficient condition. In the remainder of this paper, we will refer to this model as the gamma DPMM and assume a bivariate normal G_0 on the log-scale of the gamma kernel parameters.

3.2 *Prior specification*

When it comes to prior specification often there is not much prior knowledge on the behavior of the population of interest, but typically the researcher will have at least somewhat of an idea of

the range and midpoint/midrange of the population. We would want to set our priors to have a prior predictive distribution that encompasses this range. One way to favor a prior predictive distribution that covers the range of the data is to imagine one relatively dispersed kernel component that is centered at the midrange with 2 standard deviations either way representing the prior range. In the data illustrations in section 4, we set the range to about 2 times the data range. We can then divide the range by 4 and square that value to get the prior variance of the data. Specifically, $(range(T)/4)^2 \approx Var(T)$. This method can be implemented when fitting a gamma DPMM. We place the following distributions on the hyperparameters: $\boldsymbol{\mu} \sim N_2(a_\mu, B_\mu)$ and $\boldsymbol{\Sigma} \sim IWish(a_\Sigma, B_\Sigma)$. Making use of the moment generating function of the bivariate normal distribution, the independence property of $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$, and the first order Taylor expansion for $exp((1/2)t_.'\boldsymbol{\Sigma}t_.)$ centered around $E((1/2)t_.'\boldsymbol{\Sigma}t_.)$, we approximate $Var(T)$ as follows:

$$\begin{aligned} Var(T) &= Var(E(T|e^\theta, e^\phi)) + E(Var(T|e^\theta, e^\phi)) \\ &= E(Var(e^{\theta-\phi}|\boldsymbol{\mu}, \boldsymbol{\Sigma})) + Var(E(e^{\theta-\phi}|\boldsymbol{\mu}, \boldsymbol{\Sigma})) + E(E(e^{\theta-2\phi}|\boldsymbol{\mu}, \boldsymbol{\Sigma})) \\ &= E(e^{t_1'\boldsymbol{\mu}})E(e^{t_1'\boldsymbol{\Sigma}t_1}) + E(e^{t_2'\boldsymbol{\mu}})E(e^{t_2'\boldsymbol{\Sigma}t_2}) - E^2(e^{t_3'\boldsymbol{\mu}})E^2(e^{t_3'\boldsymbol{\Sigma}t_3}) \\ &\approx e^{t_1'a_\mu + (1/2)t_1'B_\mu t_1} e^{\frac{(1/2)t_1'B_\Sigma t_1}{a_\Sigma - d - 1}} + e^{t_2'a_\mu + (1/2)t_2'B_\mu t_2} e^{\frac{(1/2)t_2'B_\Sigma t_2}{a_\Sigma - d - 1}} \\ &\quad - e^{2(t_3'a_\mu + (1/2)t_3'B_\mu t_3)} e^{2\frac{(1/2)t_3'B_\Sigma t_3}{a_\Sigma - d - 1}} \end{aligned} \tag{8}$$

where $t_1 = (1, -2)'$, $t_2 = (2, -2)'$, $t_3 = (1, -1)'$, and $d \times d$ is the dimension of $\boldsymbol{\Sigma}$, specifically, $d = 2$. We set $a_\Sigma = 4$, which is the smallest degrees of freedom for the inverse Wishart distribution that has finite mean. If we place priors of the form $B_\mu = ((b_\mu', 0)', (0, b_\mu'))$ and $B_\Sigma = ((b_\Sigma', 0)', (0, b_\Sigma'))$ for $b_\mu', b_\Sigma' > 0$, the expression is simplified, however, we still have four parameters to specify. One solution would be to incorporate the marginal expectation:

$$E(T) \approx e^{t_3'a_\mu + \frac{1}{2}t_3'B_\mu t_3 + \frac{1}{2}t_3'B_\Sigma t_3} = e^{(a_{\mu_1} - a_{\mu_2}) + b_\mu' + b_\Sigma'} \tag{9}$$

where upon applying our earlier assumptions, we get the last expression in (9) with $a_\mu = (a_{\mu_1}, a_{\mu_2})'$. We can further simplify by setting $b_\mu' = b_\Sigma'$ resulting in two equations with three unknowns. Next, we can allocate a percentage of the marginal expectation (9) to $exp(a_{\mu_1} - a_{\mu_2})$ and a percentage of the marginal variance (8) to $exp(a_{\mu_1} - 2a_{\mu_2})$, solving for a_{μ_1} and a_{μ_2}. Finally, we can return to (8) and solve for b_μ' and b_Σ'.

Regarding the prior for α, we consider the relationship between the number of distinct components, n^*, and the value of α. In general, the number of distinct components is large for large α and small for small α. If the data set is moderately large, $E(n^*|\alpha) \approx \alpha log\left(\frac{\alpha+n}{\alpha}\right)$ can be used to

suggest an appropriate range of α values.

This approach to prior specification is based on a small amount of prior information regarding the survival distribution. In general, we recommend studying the implied prior distribution for important survival functionals, including prior point and interval estimates for the mrl function.

3.3 *Posterior inference*

Posterior simulation is simplified by truncating the mixing distribution $G_L(\cdot) \approx G(\cdot)$. Before we introduce $\boldsymbol{\theta}_l$, the first two levels of the model are,

$$\begin{aligned} t_i|\zeta_i &\overset{ind}{\sim} K(t_i;\zeta_i),\ i=1,...,n \\ \zeta_i|\boldsymbol{p},\boldsymbol{\theta} &\overset{iid}{\sim} G_L,\ i=1,...,n \end{aligned}$$

where $\boldsymbol{p}=(p_1,...,p_L)$ are the weights corresponding to the weights, $\boldsymbol{\theta}=(\boldsymbol{\theta}_1,...,\boldsymbol{\theta}_L)$. By marginalizing over the ζ_i we obtain the finite mixture model in (14). Now we can augment the model with configuration variables $\mathbf{w}=(\mathrm{w}_1,...,\mathrm{w}_n)$ such that $\mathrm{w}_i=l$ iff $\zeta_i=\boldsymbol{\theta}_l$. Using a gamma kernel ($K$) and bivariate normal baseline (G_0) with $\boldsymbol{\theta}_l=(\theta_l,\phi_l)'$, the hierarchical model is given by,

$$\begin{aligned} t_i|\boldsymbol{\theta},\mathrm{w}_i &\overset{ind}{\sim} \Gamma(t_i;e^{\theta_{\mathrm{w}_i}},e^{\phi_{\mathrm{w}_i}}),\ i=1,...,n \\ \mathrm{w}_i|\mathbf{p} &\overset{iid}{\sim} \sum_{l=1}^{L} p_l\delta_l(\mathrm{w}_i),\ i=1,...,n \\ \mathbf{p}|\alpha &\sim f(\mathbf{p};\alpha) \\ (\theta_l,\phi_l)'|\boldsymbol{\mu},\boldsymbol{\Sigma} &\overset{iid}{\sim} N_2((\theta_l,\phi_l)';\boldsymbol{\mu},\boldsymbol{\Sigma}),\ l=1,...,L \end{aligned} \tag{10}$$

where $f(\boldsymbol{p}|\alpha)=\alpha^{L-1}p_L^{\alpha-1}(1-p_1)^{-1}(1-(p_1+p_2))^{-1}\times...\times(1-\sum_{l=1}^{L-2}p_l)^{-1}$ is a special case of the generalized Dirichlet distribution (Connor & Mosimann, 1969). Here, we use conjugate priors, $\alpha\sim\Gamma(\alpha;a_\alpha,b_\alpha(\text{rate}))$, $\boldsymbol{\mu}\sim N_2(\boldsymbol{\mu};a_\mu,B_\mu)$, and $\boldsymbol{\Sigma}\sim IWish(\boldsymbol{\Sigma};a_\Sigma,B_\Sigma)$.

Now, we can utilize a blocked Gibbs sampler (Ishwaran & James, 2001) to obtain samples from the posterior distribution $p(\boldsymbol{\theta},\mathbf{w},\boldsymbol{p},\boldsymbol{\psi},\alpha|data)$ where $\boldsymbol{\psi}=(a_\mu,B_\mu,B_\Sigma)$. We have Gibbs steps for all parameters except $\boldsymbol{\theta}$, for which we use a Metropolis-Hastings step. The specifics of the posterior sampling method for the gamma DPMM are provided in Appendix B.

The posterior samples for $G_L\equiv(\boldsymbol{p},\boldsymbol{\theta})$ can be used to obtain inference for the density, survival, and hazard functions at any time point t, by directly evaluating the expressions for these functions under the gamma DPMM. Obtaining the mrl function must be done by numerical integration approximation for the integral over the survival function. From (4) we know that the mrl function at 0 returns the expected survival time, $m(0)=\mu$. Hence, the mrl function can be written alternatively

as follows:

$$m(t) \quad = \quad \frac{\int_t^\infty S(u)du}{S(t)} = \frac{\int_0^\infty S(u)du - \int_0^t S(u)du}{S(t)} = \frac{\mu - \int_0^t S(u)du}{S(t)} \tag{11}$$

We can avoid having to truncate the upper bound of the integration in the numerator in (1) by using the form of the mrl function as described in (11). We obtain posterior point and interval estimates for the mrl function by evaluating expression (11) at the posterior samples from the MCMC. We do this over a grid of survival times, $t_{0,j}$ for $j = 1, ..., m$. The survival function is monotone decreasing so the trapezoid technique is an appropriate method of approximating the integral in the mrl. We evaluate the mrl at the first grid point by $m(t_{0,1}; G_L) = [E(T; G_L) - 0.5(t_{0,1}(1 + S(t_{0,1}; G_L))]/S(t_{0,1}; G_L)$ and use the following expression for $j = 2, ..., m$:

$$m(t_{0,j}; G_L) = \frac{E(T; G_L) - \frac{1}{2}\left(t_{0,1}(1 + S(t_{0,1}; G_L)) + \sum_{i=2}^{j}(t_{0,j} - t_{0,j-1})(S(t_{0,j}; G_L) + S(t_{0,j-1}; G_L))\right)}{S(t_{0,j}; G_L)}$$

We save a lower and upper quantile along with the median at each grid point for each mixture functional to obtain (point-wise) posterior point and interval estimates.

4 DATA EXAMLES

In Section 4.1, we use simulated data to illustrate the ability of the gamma DPMM to capture non-standard mrl function shapes as well as the correlation between kernel parameters θ and ϕ. In Section 4.2, we fit a gamma DPMM as well as an exponentiated Weibull model to a data set involving survival times for subjects from two groups, including formal model comparison between the two models. In Section 4.3, we provide results of fitting the gamma DPMM to a data set of two groups both containing right censored survival times.

4.1 *Simulation examples*

In this section, we will work with two simulated data sets. The first data set consists of 200 simulations from a mixture distribution of four gamma components in which the shape and scale parameters are positively associated: $T_1 \sim 0.35\Gamma(10, 0.5) + 0.4\Gamma(20, 1) + 0.15\Gamma(30, 5) + 0.1\Gamma(40, 8)$. We fit a gamma DPMM with priors $a_\mu = (1.6, 0.4)$, $B_\mu = B_\sigma = ((0.39, 0)', (0, 0.39)')$, $a_\alpha = 2$, $b_\alpha = 1$, and $L = 40$. The effective posterior sample size is 2000. Results are shown in Figure 1.

The red dashed line represents the 95% interval estimates, the blue solid line is the point estimate, and the black dot-dashed line is the truth of the appropriate functional. We can see

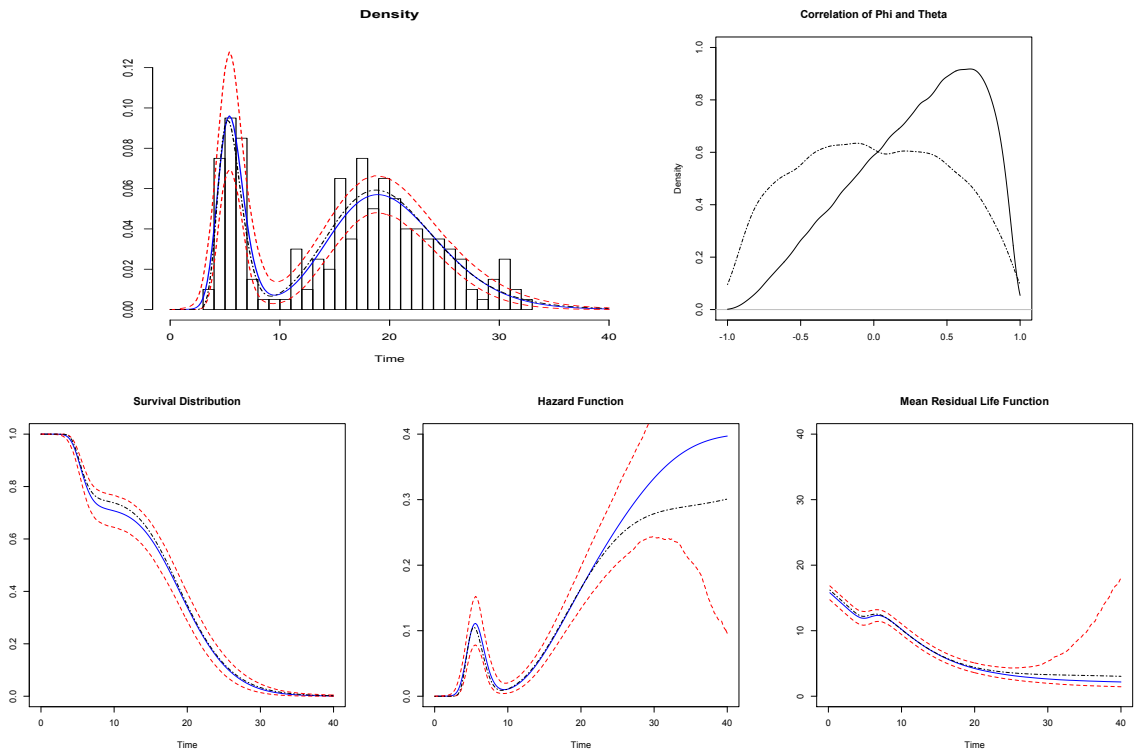

Figure 1: Simulation example 1. Posterior point (solid) and interval (dashed) estimates for the density (top left) overlaying the sample histogram and actual population density (dot-dashed), posterior (solid) and prior (dot-dashed) distribution of the correlation between θ and ϕ (top right), posterior point (solid) and interval (dashed) estimates for the survival function (lower left), hazard rate function (lower middle), and mrl function (lower right).

that truth is well within the interval estimate, moreover the point estimate is close to the truth. The correlation shows that the model (black solid) is able to capture the positive relationship between the parameters, even though the prior (black dot-dashed line) is evenly dispersed about the situation of zero correlation.

The second data set consists of 100 simulations from a distribution with negative correlation between the shape and rate parameters: $T_2 \sim 0.3\Gamma(15, 0.2)+0.25\Gamma(12, 0.5)+0.35\Gamma(8, 2)+0.1\Gamma(3, 6)$. This population was chosen to test the model's ability to separate modes that are close together, as well as model a distribution with a long tail. A gamma DPMM was fit to the data with priors $a_\mu = (2.4, -1)$, $B_\mu = B_\sigma = ((0.18, 0)', (0, 0.18)')$, $a_\alpha = 2$, $b_\alpha = 1$, and $L = 40$. The effective posterior sample size is 2000. Results are shown in Figure 2. Once again the point estimates are close to the truth even in the tail where less data is available. The uncertainty band in the mrl plot has a large upper bound, which is likely an effect of the sparse data in the tail. Numerical instability in the computation of the mrl function is also likely to contribute to the large upper

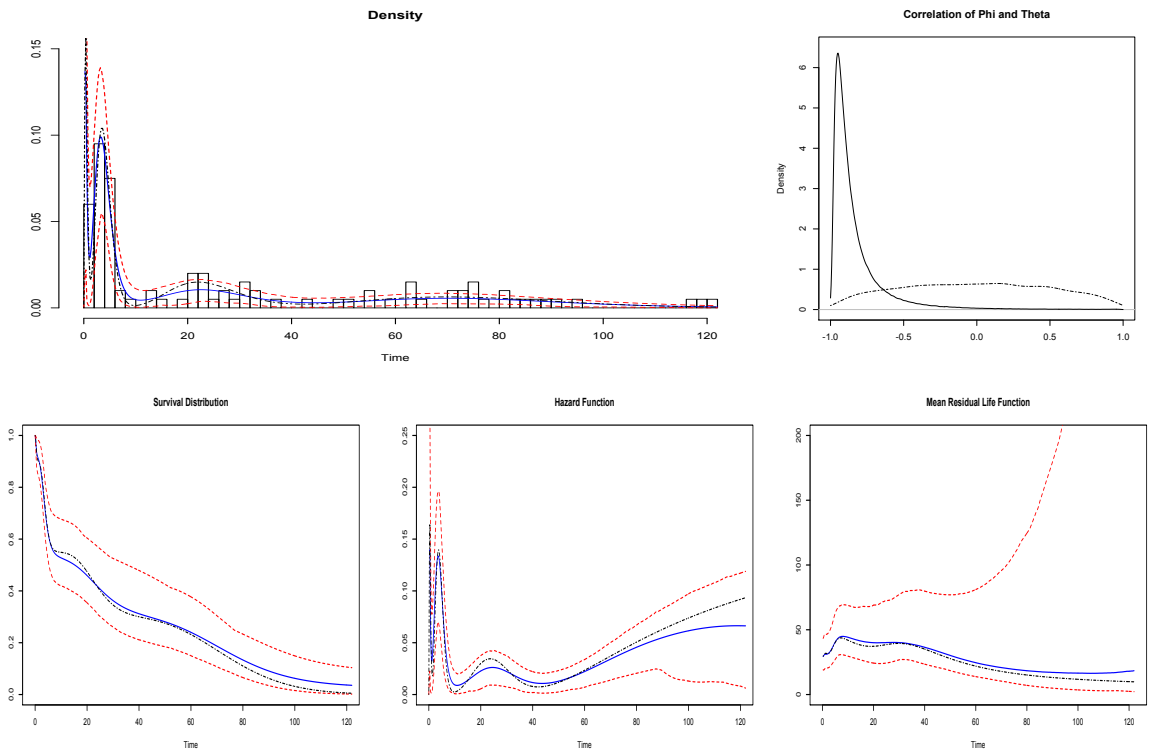

Figure 2: Simulation example 2. Posterior point (solid) and interval (dashed) estimates for the density (top left) overlaying the sample histogram and actual population density (dot-dashed), posterior (solid) and prior (dot-dashed) distribution of the correlation between θ and ϕ (top right), posterior point (solid) and interval (dashed) estimates for the survival function (lower left), hazard rate function (lower middle), and mrl function (lower right).

bound. The model picks up a strong negative correlation between the parameters.

4.2 *Analysis of survival times of rats (ad libitum vs restricted eating)*

This data set, considered earlier in Berger et al. (1988), is used to illustrate posterior inference under both an exponentiated Weibull model and the gamma DPMM. The data consists of survival times of rats in two experimental groups. The first group (Ad libitum group) is comprised of 90 rats who were allowed to eat freely as they desired. The second group (Restricted group) is comprised of 106 rats that were placed on a restricted diet.

Under the exponentiated Weibull model, we used the $P_1 = 10\%$, $P_2 = 50\%$, and $P_3 = 90\%$ quantiles of the data to obtain a system of three equations from the distribution function: $P = [1 - exp(-(Q/\sigma)^{\alpha})]^{\theta}$ where P is the percentile and Q is the survival time representing that quantile. The system of equations is solved to obtain prior means for α, σ and θ. For simplicity, exponential priors were placed on these parameters. The restricted group had respective quantile values of

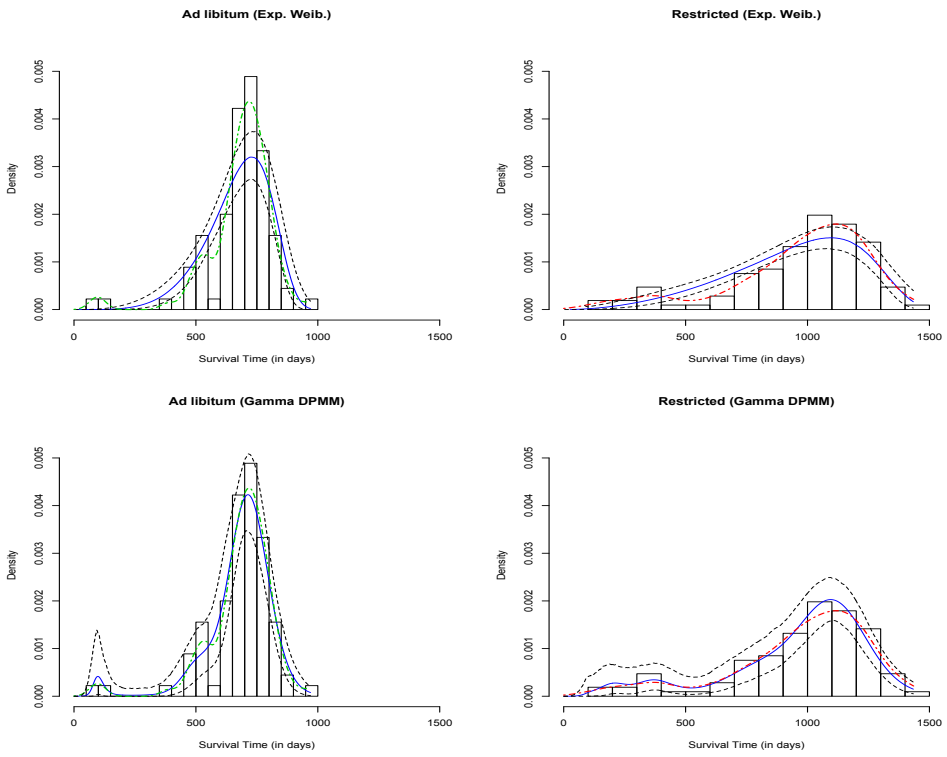

Figure 3: Relative frequency histogram and densities of lifetime (in days) of the two experimental groups (Ad libitum is left and Restricted is right) along with posterior mean and 95% interval estimates for the density functions under the exponentiated Weibull model (top) and the gamma DPMM (bottom).

($Q_1 = 1.55$, $Q_2 = 2.84$, $Q_3 = 3.34$). If we set $\alpha = 2$, $\theta = 5$, and $\sigma = 2$, then the corresponding quantiles are given as $Q_1' = 1.99$, $Q_2' = 2.85$, and $Q_3' = 4.07$ which we considered to be reasonably close to the observed quantiles. Therefore, we set hyper-parameters as $a_\alpha = 2$, $a_\theta = 5$, and $a_\sigma = 2$. Following the same methodology for the ad libitum group, we set the hyper-parameters as $a_\alpha = 4$, $a_\theta = 1$, and $a_\sigma = 2$. Posterior results were obtained using a Metropolis-Hastings algorithm in the MCMC with a trivariate normal proposal distribution on the log-scale. Point and interval estimates of the density function are plotted in the top row of Figure 3.

Prior specification for the gamma DPMM were determined using methods described in section 3.2 by allocating 60% of the marginal mean to $exp(a_{\mu_1} - a_{\mu_2})$ and 2.5% of the marginal variance to $exp(a_{\mu_1} - 2a_{\mu_2})$. For the restricted group, we use $a_\mu = (4.1, 3.6)$, $B_\mu = B_\sigma = ((0.1, 0)', (0, 0.1)')$, $a_\alpha = 2$, $b_\alpha = 1$, and $L = 40$. For the ad libitum group, we use $a_\mu = (4.16, 3.8)$, $B_\mu = B_\sigma = ((0.095, 0)', (0, 0.095)')$, $a_\alpha = 2$, $b_\alpha = 1$, and $L = 40$. The effective posterior sample size under both models are 2000 for each group. Posterior estimates for the densities for the two groups under the DP mixture model are shown in the bottom row of Figure 3. In Figure 3 we note that the

parametric model has some trouble capturing some of the characteristics of the data. In the ad libitum group (upper left) a minor mode is suggested just below the 200^{th} day. The unimodality of the exponentiated Weibull distribution makes it impossible for the parametric model to capture this shape. We note that the model tries to by reaching the tail of the estimated density out to these values, but this is at a cost of underestimating the density where most of the data exist, and overestimating the density where there is no data at all. There are many regions where the data and the density of the data (green dot-dashed) do not even fall within the interval estimates (black dashed). If we compare to how well the nonparametric model (lower left) performs we see quite a bit of improvement. The extra structure at the lower survival times is now being captured without the consequences of modeling poorly in other regions of the data. The data density remains within the interval estimates over the entire range of the data. We see similar results for the restricted group, which has a large left skew with a slight mode in the far tail. The exponentiated Weibull model (upper right) is able to model some of the skewness, but again runs into trouble by smoothing over obvious peaks and valleys. Again there are a number of regions in which the density of the data (red dot dashed) is not contained in the interval estimates of the model. The gamma DPMM (lower right) is able to capture the peaks and valleys that the exponentiated Weibull model could not. There is a slight discrepancy from the point estimate (blue solid) and the density of the data around 1250 days. Nonetheless, the data density remains within the interval estimates of the model.

By comparing the densities under the two models, there is clear evidence that the nonparametric gamma DPMM is superior to the exponentiated Weibull model. Therefore, we will use the results under the nonparametric gamma DPMM to compare the mrl functions under the two groups. In Figure 4, we plot point and interval estimates of the posterior density functions (upper left), survival functions (upper right), hazard functions (lower left), and mrl functions (lower right) for both the ad libitum (green) and restricted (red) groups. Note that the interval estimates for the mrl function are 80% probability bands as opposed to the other interval estimates, which are 95% probability bands. The reason for this is to reduce the steepness for which the upper bound shoots upward towards the tail of the data. This is likely due to the lower number of observations towards the end of the range of the data and also the numerical instability in computing the mrl function. Looking at the estimated densities survival functions we can see that the majority of the ad libitum group have lower survival times compared to the restricted group. The mrl functions are monotonically decreasing and do not cross with regard to the point estimates. Moreover the interval estimates do not cross until the we reach about 800 days. This leads us to conclude that the remaining life expectancy of a rat in the restricted group is higher than the remaining life expectancy of a rat in the ad libitum group until we reach about 800 days.

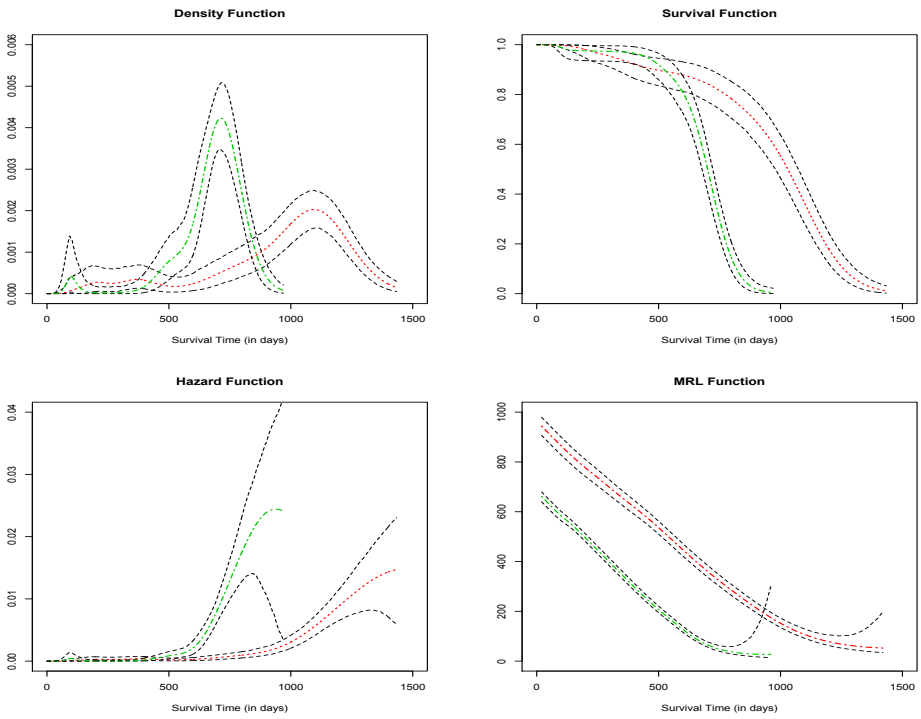

Figure 4: Point and interval estimates of lifetime (in days) for the density (top left), survival (top right), hazard rate (lower left), and ml (lower right) functions of the two experimental groups under the gamma DPMM.

We use the minimum posterior predictive loss approach (Gelfand & Ghosh, 1998) to compare the exponentiated Weibull model to the nonparametric gamma DPMM. Under this criterion the goal is to minimize, within the collection of models under consideration, the expectation of a specified loss function under the posterior predictive distribution of replicate responses $\boldsymbol{t}_{rep}$ given the observed data $\boldsymbol{t}_{obs}$. Here, we use the square error loss function so that the general criterion is given by $D_k(m) = \sum_{i=1}^{n} var(t_{i,rep}|\boldsymbol{t}_{obs}, m) + \frac{k}{k+1}\sum_{i=1}^{n}(E(t_{i,rep}|\boldsymbol{t}_{obs}, m) - t_{i,obs})^2$, where $t_{i,rep}$ is a replicate of the i^{th} observation, $t_{i,obs}$, under the posterior predictive distribution of the m^{th} model. The first term is representative of a penalty measure $P(m)$, and the second term is a goodness-of-fit measure $G(m)$. The value of k is specified as the relative regret for departure from $t_{i,rep}$. Note that as k tends to infinity, the criterion becomes the sum of the penalty and goodness-of-fit terms.

For the exponentiated Weibull model (m_1), obtaining $E(t_{i,rep}|\boldsymbol{t}_{obs}, m)$ and $var(t_{i,rep}|\boldsymbol{t}_{obs}, m)$ is straightforward. The posterior predictive distribution is given by $p(t_{i,rep}|\boldsymbol{t}_{obs}) = \int EW(t_{i,rep}|\alpha, \theta, \sigma)$ $\times\, p(\alpha, \theta, \sigma|data)d\alpha d\theta d\sigma$ and can thus be sampled by taking the posterior samples $(\alpha_b, \theta_b, \sigma_b)$, for $b = 1, ..., B$, and drawing $t_{i,rep,b}$ from the exponentiated Weibull distribution given each posterior parameter vector. Next, we compute the mean and variance of the B replicates. Important to

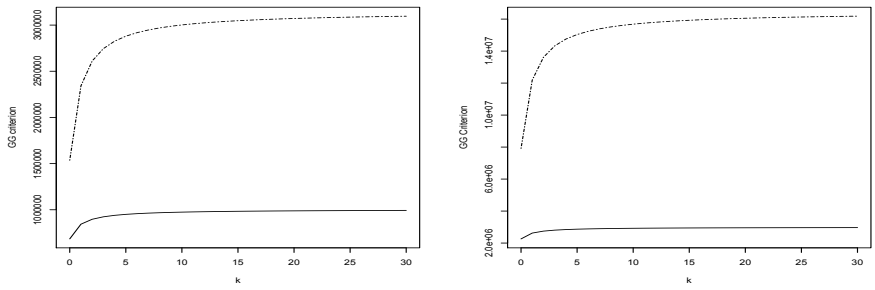

Figure 5: Values of the posterior predictive loss criterion for comparison between the parametric exponentiated Weibull model (dot dashed lines) and nonparametric gamma DPMM (solid lines).

note is that the mean and variance for one experimental group is going to be the same for each observation in that group. We find the $E(t_{i,rep}|\boldsymbol{t}_{obs}, m_1)$ and $var(t_{i,rep}|\boldsymbol{t}_{obs}, m_1)$ for the ad libitum group to be 671.2 and 17433.0, respectively, and for the the restricted group to be 949.5 and 74691.7, respectively. Thus the ad libitum group has $G(m_1)^a = \sum_{i=1}^{90}(671.2 - t_{i,obs})^2 = 1615787$ and $P(m_1)^a = 90*(17433.0) = 1568967$. The restricted group has $G(m_1)^r = \sum_{i=1}^{106}(949.5-t_{i,obs})^2 = 8542725$ and $P(m_1)^r = 106 * (74691.7) = 7917319$.

Obtaining the criterion under the nonparametric gamma DPMM (m_2) takes a little more care. Recall that $t_i|G \overset{ind}{\sim} \int \Gamma(t_i; exp(\theta), exp(\phi))dG(\theta, \phi)$ for $i = 1, ..., n$. In order to obtain replicates for each t_i, we need to know the l^{th} component from which the observed t_i came from according to the model, ie. $t_{i,rep}|\boldsymbol{t}_{obs}, m_2 \sim \int \Gamma(t_{i,rep}|exp(\theta_{l_i}), exp(\phi_{l_i}))p(exp(\theta_{l_i}), exp(\phi_{l_i})|data)d\theta_{l_i}d\phi_{l_i}$, for $i = 1, ..., n$, where the subscript l_i is the i^{th} value of the posterior sample of $\mathbf{w}$ and θ_{l_i} and ϕ_{l_i} are the l_i^{th} posterior samples of θ and ϕ. Essentially a single $t_{i,rep}$ is sampled from the gamma distribution at each posterior iteration $b = 1, ..., B$ integrating out all possible values of θ_{l_i} and ϕ_{l_i}. After obtaining B $t_{i,rep}$, we compute the mean ($E(t_{i,rep}|\boldsymbol{t}_{obs}, m_2)$) and variance ($var(t_{i,rep}|\boldsymbol{t}_{obs}, m_2)$) at each i^{th} replicate. For the ad libitum group we obtained $G(m_2)^a = 318919.2$ and $P(m_2)^a = 684342.1$, and for the restricted group $G(m_2)^r = 739435$ and $P(m_2)^r = 2247120$. Figure 5 is a plot of the criterion values over a grid of k values. For both groups the nonparametric gamma DPMM performs significantly better than the exponentiated Weibull model. The results of the formal model comparison support our earlier argument that the nonparametric gamma DPMM is indeed a better model for these data.

4.3 *Analysis of survival times of patients with small cell lung cancer*

As an example of obtaining inference in the presence of right censoring, we fit a gamma DPMM to the survival times, in days, of two treatment groups of patients with small cell lung cancer Ying

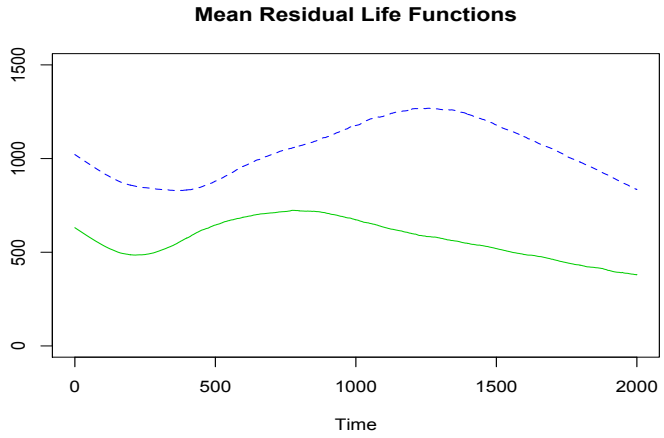

Figure 6: Posterior point estimates for the mrl functions of Arm A (blue dashed) and Arm B (green solid).

et al. (1995). The patients were randomly assigned to one of two treatments referred to as Arm A and Arm B. Arm A patients received cisplatin (P) followed by etoposide (E), while Arm B patients received (E) followed by (P). There were a total of 62 patients in Arm A with 15 right censored survival times, while Arm B consisted of 59 patients with 8 right censored survival times. We fit a gamma DPMM independently to the two groups. We allocated 60% of the mean to $exp(a_{\mu_1} - a_{\mu_2})$ and 2.25% to $exp(a_{\mu_1} - 2a_{\mu_2})$ resulting in the following priors for Arm A: $a_\mu = (2.5, -3)$, $B_\mu = B_\sigma = ((0.21, 0)', (0, 0.21)')$. Analoguously, for Arm B: $a_\mu = (2.6, -2.9)$, $B_\mu = B_\sigma = ((0.21, 0)', (0, 0.21)')$. We use $a_\alpha = 3$, $b_\alpha = 1$, and $N = 25$. The effective posterior sample size is 2000.

The point estimates for mrl functions of the two treatment groups show Arm A to have a consistently higher mean residual life compared to Arm B in Figure 6. The result leads us to believe that Arm A treatment is more effective than the Arm B treatment. We take a closer look at the difference of the mean residual life survival times at a number of fixed time points. Specifically, in Figure 7, we explore the posterior density of $m_A(t) - m_B(t)$, where $m_A(t)$ and $m_B(t)$ are the mrl functions corresponding to Arm A and Arm B, respectively, at six time points. Time zero, which is the estimated difference of the overall mean of the two distributions, depicts a clear difference between the two treatments in favor of Arm A. The same is true at 100 days, and to a somewhat smaller extent, at 250 days. At the larger time points, although the indication of a difference is still present, it becomes less emphatic.

A further means of comparing the mrl functions of Arm A and Arm B is through the probability of the mrl function of one group being higher than that of the other over a grid of survival times. Figure 8 plots as a function of time the prior probability, $\Pr(m_A(t) > m_B(t))$, and the posterior probability, $\Pr(m_A(t) > m_B(t)|data)$. The prior probability is relatively flat around 0.5, and hence

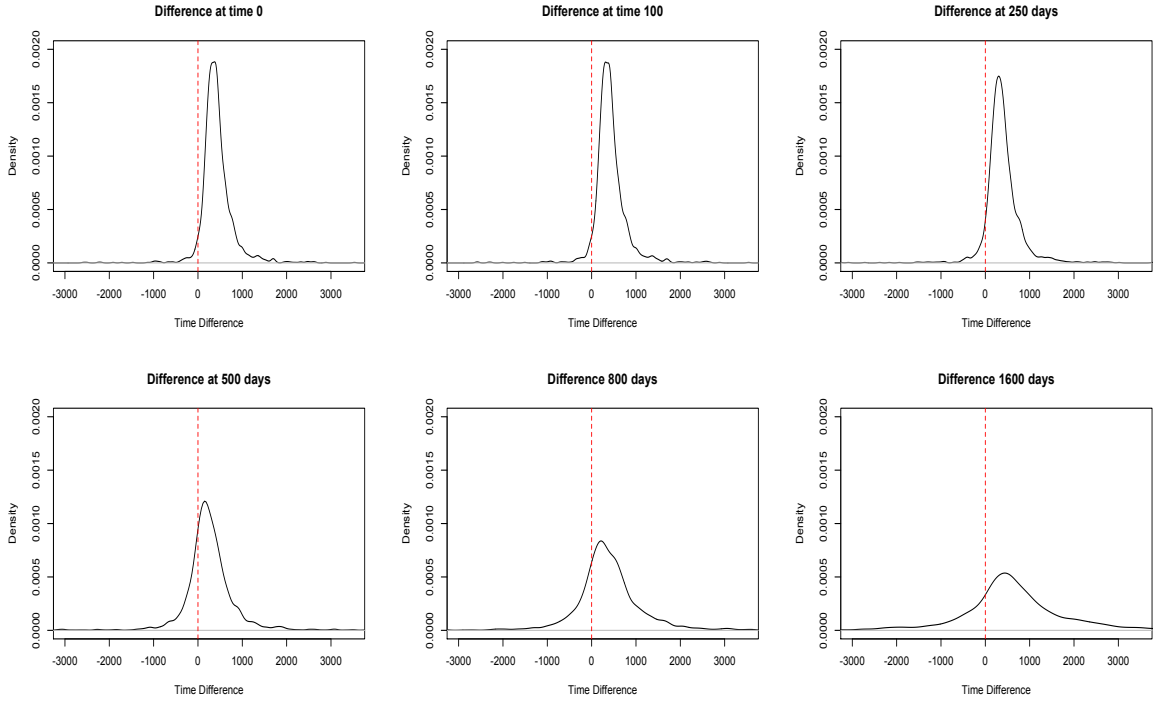

Figure 7: Densities of the mrl of Arm A minus Arm B at a number of fixed time points.

the prior specification does not favor either of the two groups. The posterior probabilities strongly suggest Arm A has the higher mrl function. The posterior probability is particularly high during the early time period, decreases slightly around 500 days, followed by another peak around 1200 days. The posterior probability remains above 0.7 across the range of survival times in the data.

5 DISCUSSION

We discussed the benefits of Dirichlet process mixture modeling to obtain flexible inference for the mrl function. With the focus on inference for this particular functional, the choice of the mixture kernel plays an important role. Under the sufficient condition given in Section 3.1, we studied restrictions that need to be placed on the mixture model in order to ensure that the mrl function of the mixture distribution is well defined. In addition, we provided a result on the tail behavior of the mixture mrl function based on the corresponding property for the mrl function of the kernel distribution. The gamma kernel was shown to possess the most desirable properties among the distributions we investigated. We showed that under a gamma mixture, the resulting mrl function is dense in the pointwise sense on the space of continuous mrl functions. The practical utility of the proposed nonparametric mixture model was demonstrated through analysis of simulated data examples and real data sets from the literature.

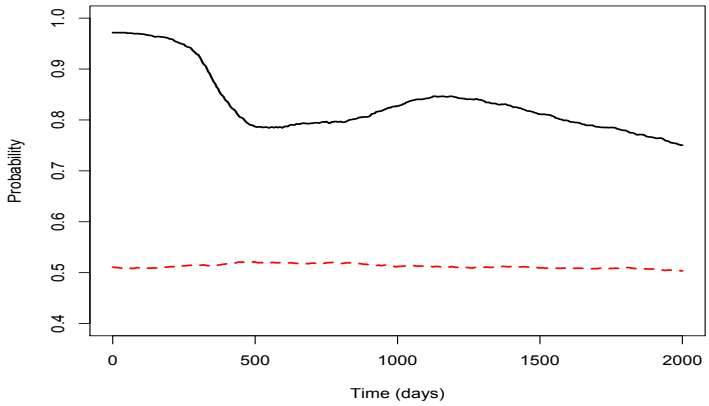

Figure 8: The posterior (solid line) and prior (dashed line) probability of the mrl function of Arm A being higher than the mrl function of Arm B, as a function of time.

A practically important extension involves methods for inclusion of covariates. Mean residual life regression modeling can be explored under either a structured semiparametric setting, such as the proportional mrl setting, or a fully nonparametric framework, for instance, based on mixture modeling for the joint response-covariate distribution. Moreover, of both methodological and practical interest is development of nonparametric prior models directly for the mrl function. The support of the nonparametric prior would have to be functions that satisfy the properties in the characterization theorem of the mrl function. Under such a prior, inference for the entire distribution would be obtainable by using the inversion formula.

References

Abdous, B. & Berred, A. (2005). Mean residual life estimation. *Statistical Planning and Inference* **132**, 3–19.

Berger, R. L., Boos, D. D. & Guess, F. M. (1988). Tests and confidence sets for comparing two mean residual life functions. *Biometrics* **44**, 103–115.

Chen, Y. Q. & Cheng, S. (2005). Semiparametric regression analysis of mean residual life with censored data. *Biometrika* **92**, 19–29.

Connor, R. J. & Mosimann, J. E. (1969). Concepts of independence for proportions with a generalization of the Dirichlet distribution. *Journal of the American Statistical Association* **64**, 194–206.

Ferguson, T. S. (1973). A Bayesian analysis of some nonparametric problems. *The Annals of Statistics* **1**, 209–230.

Finkelstein, M. S. (2002). On the shape of the mean residual lifetime function. *Applied Stochastic Models in Business and Industry* **18**, 135–146.

Gelfand, A. E. & Ghosh, S. K. (1998). Model choice: A minimum posterior predictive loss approach. *Biometrika* **85**.

Govilt, K. & Aggarwal, K. (1983). Mean residual life function of normal, gamma and lognormal densities. *Reliability Engineering* **5**, 47–51.

Guess, F. & Proschan, F. (1985). Mean residual life: theory and applications. Technical Report 85-178, North Carolina State University and Florida State University, Tallahassee,Florida.

Gupta, R. C. & Akman, H. O. (1995). Mean residual life functions for certain types of non-monotonic ageing. *Commun. Statist.-Stochastic Models* **11**, 219–225.

Gupta, R. C., Akman, O. & Lvin, S. (1999). A study of log-logistic model in survival analysis. *Biometrical Journal* **41**, 431–443.

Hall, W. J. & Wellner, J. A. (1979). Estimation of mean residual life. University of Rochester.

Hall, W. J. & Wellner, J. A. (1981). Mean residual life. In *Statistics and Related Topics*, Eds. M. Csörgö, D. Dawson, J. Rao & A. Saleh. North-Holland Publishing Company.

Hanson, T. E. (2006). Modeling censored lifetime data using a mixture of gammas baseline. *Bayesian Analysis* **1**, 575–594.

Ishwaran, H. & James, L. F. (2001). Gibbs sampling methods for stick-breaking priors. *American Statistical Association* **96**, 161–173.

Johnson, M. A. & Taaffe, M. R. (1998). The denseness of Phase distributions. Research Memorandum No 88-20.

Johnson, W. O. (1999). Survival analysis for interval data. In *Institute for Mathematics and Its Applications: Statistics in the Health Sciences: Diagnosis and Prediction*, Eds. P. Grambsch & S. Geisser, volume 114, pp. 75–90. Springer-Verlag.

Kottas, A. (2006). Nonparametric Bayesian survival analysis using mixtures of Weibull distributions. *Journal of Statistical Planning and Inference* **136**, 578–596.

Kuo, L. & Mallick, B. (1997). Bayesian semiparametric inference for the accelerated failure-time model. *The Canadian Journal of Statistics* **25**, 457–472.

Lahiri, P. & Park, D. H. (1991). Nonparametric Bayes and empirical Bayes estimators of mean residual life. *Journal of Statistical Planning and Inference* **29**, 125–136.

Maguluri, G. & Zhang, C.-H. (1994). Estimation in the mean residual life regression model. *Journal of the Royal Statistical Society Series B* **56**, 477–489.

Mudholkar, G. S. & Strivasta, D. K. (1993). Exponentiated Weibull family for analyzing bathtub failure-rate data. *IEEE Transactions of Reliability* **42**, 299–302.

Oakes, D. & Dasu, T. (1990). A note on mean residual life. *Biometrika* **77**, 409–410.

Pham, H. & Lai, C.-D. (2007). On recent generalizations of the Weibull distribution. *IEEE Transactions on Reliability* **56**, 454–458.

Sethuraman, J. (1994). A constructive definition of Dirichlet priors. *Statistica Sinica* **4**, 639–650.

Smith, P. J. (2002). Analysis of failure and survival data. In *Texts in Statistical Science Series*, Eds. C. Chatfield, J. Lindsey, M. Tanner & J. Zidek, pp. –. Chapman and Hall/CRC.

Xie, M., Goh, T. & Tang, Y. (2004). On changing points of mean residual life and failure rate function for some generalized Weibull distributions. *Reliability Engineering and System Safety* **84**, 293–299.

Yang, G. L. (1978). Estimation of a biometric function. *The Annals of Statistics* **6**, 112–116.

Ying, Z., Jung, S. & Wei, L. (1995). Survival analysis with median regression models. *Journal of the American Statistical Association* **90**, 178–184.

APPENDIX A: Proof of the Lemmas

A.1: Proof of Lemma 1

Let $f(\cdot)$ and $S(\cdot)$ be the kernel density and survival functions, respectively, of a DPMM. Assume that $f\cdot) > 0$ for all t$\geq$ 0 or for all $t > t_0$ where $t_0 \geq 0$ is some finite value. The corresponding mrl function of the DPMM with L components is given by:

$$m(t; G_L) = \frac{\int_t^\infty \sum_{l=1}^L p_l S(u;\boldsymbol{\theta}_l)du}{\sum_{l=1}^L p_l S(t;\boldsymbol{\theta}_l)} = \frac{\sum_{l=1}^L p_l \int_t^\infty S(u;\boldsymbol{\theta}_l)du}{\sum_{l=1}^L p_l S(t;\boldsymbol{\theta}_l)}$$

Note that $\lim_{t\to\infty}\sum_{l=1}^{L} p_l \int_t^{\infty} S(u;\boldsymbol{\theta}_l)du = 0$ and $\lim_{t\to\infty}\sum_{l=1}^{L} p_l S(t;\boldsymbol{\theta}_l) = 0$, so by L'Hopital's Rule we have:

$$\lim_{t\to\infty} m(t;G_L) = \lim_{t\to\infty} \frac{\frac{d}{dt}\sum_{l=1}^{L} p_l \int_t^{\infty} S(u;\boldsymbol{\theta}_l)du}{\frac{d}{dt}\sum_{l=1}^{L} p_l S(t;\boldsymbol{\theta}_l)} = \lim_{t\to\infty} \frac{-\sum_{l=1}^{L} p_l S(t;\boldsymbol{\theta}_l)}{-\sum_{l=1}^{L} p_l f(t;\boldsymbol{\theta}_l)}$$

Once again the limit as t goes to infinity of both the numerator and denominator is zero, so applying L'Hopital once more we have the following:

$$\lim_{t\to\infty} m(t;G_L) = \lim_{t\to\infty} \frac{-\sum_{l=1}^{L} p_l f(t;\boldsymbol{\theta}_l)}{\sum_{l=1}^{L} p_l f'(t;\boldsymbol{\theta}_l)}$$

Suppose that the mrl function, $m(t;\boldsymbol{\theta})$ of the kernel distribution tends to infinity as $t \to \infty$. Then,

$$\lim_{t\to\infty} m(t;\boldsymbol{\theta}) = \lim_{t\to\infty} \frac{\int_t^{\infty} S(u;\boldsymbol{\theta})du}{S(t;\boldsymbol{\theta})} = \lim_{t\to\infty} \frac{-S(t;\boldsymbol{\theta})}{-f(t;\boldsymbol{\theta})} = \lim_{t\to\infty} \frac{-f(t;\boldsymbol{\theta}_l)}{f'(t;\boldsymbol{\theta})} = \infty$$

In other words, $lim_{t\to\infty}(-f(t;\boldsymbol{\theta})/f'(t;\boldsymbol{\theta})) = 0$, so $-f(x;\boldsymbol{\theta})$ grows at a much faster rate the $f'(t;\boldsymbol{\theta})$. Hence $f'(t;\boldsymbol{\theta})$ is " little-o " of $-f(t;\boldsymbol{\theta})$: $f'(t;\boldsymbol{\theta}) \in o(-f(t;\boldsymbol{\theta}))$. Returning now to the mrl function of the DPMM, for each component $l = 1, ..., L$: $f'(t;\boldsymbol{\theta}_l) \in o(-f(t;\boldsymbol{\theta}_l)) \Leftrightarrow p_l f'(t;\boldsymbol{\theta}_l) \in o(-p_l f(t;\boldsymbol{\theta}_l))$ since by definition of "little - o" $\forall\epsilon > 0$ there exists $t_{0l} \in \mathbb{R}$ such that $|f'(t;\boldsymbol{\theta}_l)| \leq \epsilon|-f(t;\boldsymbol{\theta}_l)|\forall t \geq t_{0l}$ for $l = 1, ...L$. Multiplying by $p_l > 0$ on either side gives $p_l|f'(t;\boldsymbol{\theta}_l)| \leq \epsilon p_l|-f(t;\boldsymbol{\theta}_l)| \Leftrightarrow |p_l f'(t;\boldsymbol{\theta}_l)| \leq \epsilon| - p_l f(t;\boldsymbol{\theta}_l)|$. Since this last inequality holds for all $l = 1, ..., L$ we can apply the sum over l on both sides obtaining $\sum_{l=1}^{L} |p_l f'(t;\boldsymbol{\theta}_l)| \leq \sum_{l=1}^{L} \epsilon| - p_l f(t;\boldsymbol{\theta}_l)|$. We can bound the left side of the inequality below using the triangle inequality, $|\sum_{l=1}^{L} p_l f'(t;\boldsymbol{\theta}_l)| \leq \sum_{l=1}^{L} |p_l f'(t;\boldsymbol{\theta}_l)|$. Meanwhile, the right side of the inequality can be written as $\sum_{l=1}^{L} \epsilon| - p_l f(t;\boldsymbol{\theta}_l)| = \epsilon| - \sum_{l=1}^{L} p_l f(t;\boldsymbol{\theta}_l)|$ since $-p_l f(t;\boldsymbol{\theta}_l) \leq 0$ for each $l = 1, ..., L$. Thus, we can make the following statement: $\forall\epsilon > 0 \exists t_0 \in \mathbb{R}$ such that $|\sum_{l=1}^{L} p_l f'(t;\boldsymbol{\theta}_l)| \leq \epsilon| - \sum_{l=1}^{L} p_l f(t;\boldsymbol{\theta}_l)|\forall t \geq max\{t_{01}, t_{02}, ..., t_{0L}\}$. In other words, $\sum_{l=1}^{L} p_l f'(t;\boldsymbol{\theta}_l) \in o(-\sum_{l=1}^{L} p_l f(t;\boldsymbol{\theta}_l))$, and therefore,

$$\lim_{t\to\infty} m(x;G_L) = \lim_{t\to\infty} \frac{-\sum_{l=1}^{L} p_l f(t;\boldsymbol{\theta}_l)}{\sum_{l=1}^{L} p_l f'(t;\boldsymbol{\theta}_l)} = \infty$$

Now suppose that the mrl function of the kernel distribution tends to zero as $t \to \infty$. Hence, we can say that $S(t;\boldsymbol{\theta}) \in o(f(t;\boldsymbol{\theta}))$, since we have the following:

$$\lim_{t\to\infty} m(t;G_L) = \lim_{t\to\infty} \frac{S(t;\boldsymbol{\theta})}{f(t;\boldsymbol{\theta})} = 0$$

Thus, in the DPMM we have for each component, $l = 1, ..., L$, $S(t;\boldsymbol{\theta}_l) \in o(f(t;\boldsymbol{\theta}_l)) \Leftrightarrow p_l S(t;\boldsymbol{\theta}_l) \in o(p_l f(t;\boldsymbol{\theta}_l))$. From the definition $\forall\epsilon > 0 \exists t_{0l} \in \mathbb{R}$ such that $|p_l S(t;\boldsymbol{\theta}_l)| \leq \epsilon|p_l f(t;\boldsymbol{\theta}_l))|\forall t \geq t_{0l}$. Applying the sum over the components gives us $\sum_{l=1}^{L} |p_l S(t;\boldsymbol{\theta}_l)| \leq \sum_{l=1}^{L} \epsilon|p_l f(t;\boldsymbol{\theta}_l))|$. The left side of the inequality can be written as $\sum_{l=1}^{L} |p_l S(t;\boldsymbol{\theta}_l)| = |\sum_{l=1}^{L} p_l S(t;\boldsymbol{\theta}_l)|$, similarly, the right side can be written as $\sum_{l=1}^{L} \epsilon|p_l f(t;\boldsymbol{\theta}_l))| = \epsilon|\sum_{l=1}^{L} p_l f(t;\boldsymbol{\theta}_l))|$. Hence, $\forall\epsilon > 0 \exists t_0 \in \mathbb{R}$ such that $|\sum_{l=1}^{L} p_l S(t;\boldsymbol{\theta}_l)| \leq \epsilon|\sum_{l=1}^{L} p_l f(t;\boldsymbol{\theta}_l))|\forall t \geq max\{t_{01}, t_{02}, ..., t_{0L}\}$. In other words, $\sum_{l=1}^{L} p_l S(t;\boldsymbol{\theta}_l) \in o(\sum_{l=1}^{L} p_l f(t;\boldsymbol{\theta}_l))$, and therefore,

$$\lim_{t\to\infty} m(t;G_L) = lim_{t\to\infty}\frac{\sum_{l=1}^{L} p_l S(t;\boldsymbol{\theta}_l)}{\sum_{l=1}^{L} p_l f(t;\boldsymbol{\theta}_l)} = 0$$

A.2: Proof of Lemma 2

Let $\mathscr{F}$ be the space of absolutely continuous distribution functions on $\mathbb{R}^+$ with finite mean, $\mu < \infty$. Let $\mathscr{M}$ be the space of continuous mrl functions. Consider the class of gamma mixture distributions, $\mathcal{C}$. Now, let $m(t)$, for $t \geq 0$, be any mrl function in $\mathscr{M}$. We can obtain the survival function corresponding to $m(t)$ via the Inversion Formula:

$$S(t) = \frac{m(0)}{m(t)} exp\left(-\int_0^t \frac{1}{m(u)} du\right)$$

Hence the corresponding distribution function is defined by $F(t) = 1 - S(t)$. Now, we know that $\mathcal{C}$ is dense in $\mathscr{F}$. Particularly, if we define a sequence of distribution functions, $\{F_n(t)\} \subseteq \mathcal{C}$, as follows:

$$F_n(t) = \sum_{l=1}^{\infty} \left[F\left(\frac{l}{n}\right) - F\left(\frac{l-1}{n}\right)\right] F_\Gamma(t; l, n)$$

where $\left[F\left(\frac{l}{n}\right) - F\left(\frac{l-1}{n}\right)\right]$ are the corresponding weights of the gamma cumulative distributions functions, $F_\Gamma(t; l, n)$, with shape parameter l and rate parameter n.. Johnson & Taaffe (1998) show that for any $t_0 \geq 0$, $\lim_{n\to\infty} F_n(t_0) = F(t_0)$. That is the sequence $\{F_n(t)\}$ converges weakly (or pointwise) to $F(t)$. For the case of a finite mixture, the sequence is defined such that the limit of the sequence as the number of mixture components tends to infinity is also taken. Note that since $\{F_n(t)\}$ converges weakly (or pointwise) to $F(t)$, then the sequence of survival functions, $\{S_n(t)\}$ converges pointwise to $S(t)$, since $\lim_{n\to\infty} S_n(t) = \lim_{n\to\infty}(1 - F_n(t)) = 1 - \lim_{n\to\infty} F_n(t) = 1 - F(t) = S(t)$.

Define the sequence of mrl functions, $\{m_n(t)\}$, through the sequence of survival function, $\{S_n(t)\}$ by the following,

$$m_n(t) = \frac{\int_t^\infty S_n(u)du}{S_n(t)}$$

Consider any $t_0 \geq 0$, then take the limit of the sequence,

$$\lim_{n\to\infty} m_n(t_0) \quad = \quad \lim_{n\to\infty} \frac{\int_{t_0}^\infty S_n(u)du}{S_n(t_0)} = \frac{\lim_{n\to\infty}\int_{t_0}^\infty S_n(u)du}{\lim_{n\to\infty} S_n(t_0)}$$

The limit can be distributed in the last step as a basic property of limits provided the limits exist and the limit of the denominator is not zero. Upon evaluating these limits, we will show all these requirements are met. The bottom limit is trivial since $\{S_n(t)\}$ converges pointwise to $S(t)$ which is bounded by $0 < S(t) \leq 1$. The nontrivial step is being able to move the limit inside the integral

in the numerator. We can rewrite $\int_t^\infty S_n(u)du$ as $\mu_n - \int_0^t S_n(u)du$, where μ_n is the mean of the nth distribution in the sequence.

$$\Rightarrow \lim_{n\to\infty} m_n(t_0) = \frac{\lim_{n\to\infty}\left[\mu_n - \int_0^{t_0} S_n(u)du\right]}{\lim_{n\to\infty} S_n(t_0)} = \frac{\lim_{n\to\infty}\mu_n - \lim_{n\to\infty}\int_0^{t_0} S_n(u)du}{\lim_{n\to\infty} S_n(t_0)}$$

Now, we will establish that μ_n is finite for every n, and that $\lim_{n\to\infty}\mu_n = \mu$:

$$\begin{aligned}
\mu_n &= \int_0^\infty S_n(u)du = \int_0^\infty 1 - F_n(u)du = \int_0^\infty 1 - \sum_{l=1}^\infty \left[F\left(\frac{l}{n}\right) - F\left(\frac{l-1}{n}\right)\right] F_\Gamma(u;l,n)du \\
&= \int_0^\infty \sum_{l=1}^\infty \left[F\left(\frac{l}{n}\right) - F\left(\frac{l-1}{n}\right)\right] - \sum_{l=1}^\infty \left[F\left(\frac{l}{n}\right) - F\left(\frac{l-1}{n}\right)\right] F_\Gamma(u;l,n)du \\
&= \int_0^\infty \sum_{l=1}^\infty \left[F\left(\frac{l}{n}\right) - F\left(\frac{l-1}{n}\right)\right] (1 - F_\Gamma(u;l,n))du \\
&= \int_0^\infty \sum_{l=1}^\infty \left[F\left(\frac{l}{n}\right) - F\left(\frac{l-1}{n}\right)\right] S_\Gamma(u;l,n)du
\end{aligned}$$

where $S_\Gamma(u;l,n)$ is the survival function of the Gamma distribution with shape l and rate n. Now, since $\left[F\left(\frac{l}{n}\right) - F\left(\frac{l-1}{n}\right)\right] S_\Gamma(u;l,n) \geq 0$, by Tonelli's Theorem, we can exchange the summation and integral:

$$\begin{aligned}
\Rightarrow \mu_n &= \sum_{l=1}^\infty \left[F\left(\frac{l}{n}\right) - F\left(\frac{l-1}{n}\right)\right] \int_0^\infty S_\Gamma(u;l,n)du \;=\; \sum_{l=1}^\infty \left[F\left(\frac{l}{n}\right) - F\left(\frac{l-1}{n}\right)\right] \left(\frac{l}{n}\right) \\
&= \sum_{l=1}^\infty \left[\int_{(l-1)/n}^{l/n} f(u)\,du\right] \left(\frac{l}{n}\right) \;=\; \sum_{l=1}^\infty \int_{(l-1)/n}^{l/n} \left(\frac{l}{n}\right) f(u)\,du \\
&\leq \sum_{l=1}^\infty \int_{(l-1)/n}^{l/n} \left(u + \frac{1}{n}\right) f(u)\,du \;=\; \sum_{l=1}^\infty \int_{(l-1)/n}^{l/n} (u)\, f(u)\,du + \sum_{l=1}^\infty \int_{(l-1)/n}^{l/n} \left(\frac{1}{n}\right) f(u)\,du \\
&= \int_0^\infty uf(u)du + \frac{1}{n}\int_0^\infty f(u)du \;=\; \mu + \frac{1}{n} < \infty
\end{aligned}$$

where the last inequality holds since we are assuming F has finite mean. The inequality, $\mu_n \leq \mu + (1/n)$, also provides the following upper bound for the limit:

$$\Rightarrow \lim_{n\to\infty}\mu_n \leq \lim_{n\to\infty}\left(\mu + \frac{1}{n}\right) = \mu$$

We can also establish the following lower bound for the limit,

$$\lim_{n\to\infty}\mu_n = \lim_{n\to\infty}\sum_{l=1}^\infty \int_{(l-1)/n}^{l/n} \left(\frac{l}{n}\right) f(u)du \geq \lim_{n\to\infty}\sum_{l=1}^\infty \int_{(l-1)/n}^{l/n} uf(u)du = \lim_{n\to\infty}\int_0^\infty uf(u)du = \mu$$

Therefore, by Squeeze Theorem, $\lim_{n\to\infty}\mu_n = \mu$.

Using the fact that $S_n(t) \leq 1$ and $S_n(t)$ converges pointwise as $n \to \infty$ to $S(t)$, by the dominated convergence theorem,

$$\lim_{n\to\infty} \int_0^{t_0} S_n(u)du = \int_0^{t_0} \lim_{n\to\infty} S_n(u)du = \int_0^{t_0} S(u)du$$

Returning to the limit of the sequence of mrl functions,

$$\lim_{n\to\infty} m_n(t_0) = \frac{\lim_{n\to\infty} \mu_n - \lim_{n\to\infty} \int_0^{t_0} S_n(u)du}{\lim_{n\to\infty} S_n(t_0)} = \frac{\mu - \int_0^{t_0} S(u)du}{S(t_0)} = \frac{\int_{t_0}^{\infty} S(u)du}{S(t_0)} = m(t_0).$$

Hence, $\{m_n(t)\}$ convergence pointwise to $m(t)$, providing the denseness result for continuous mrl functions under gamma mixture distributions.

APPENDIX B: Posterior sampling from the gamma DPMM

As we stated in the text, posterior samples of the unknown parameters can easily be obtained using the block Gibbs sampler for DP mixtures described in Ishwaran & James (2001). Recall that our full hierarchical model is given by,

$$\begin{aligned}
t_i|\boldsymbol{\theta}, \mathrm{w}_i &\overset{ind}{\sim} \Gamma(t_i; e^{\theta_{\mathrm{w}_i}}, e^{\phi_{\mathrm{w}_i}}) \\
\mathrm{w}_i|\mathbf{p} &\overset{iid}{\sim} \sum_{l=1}^{L} p_l \delta_l(\mathrm{w}_i) \\
\mathbf{p}|\alpha &\sim f(\mathbf{p};\alpha) \quad (SB) \\
\boldsymbol{\theta}_l = (\theta_l, \phi_l)'|\boldsymbol{\mu}, \boldsymbol{\Sigma} &\overset{iid}{\sim} N_2((\theta_l, \phi_l)'; \boldsymbol{\mu}, \boldsymbol{\Sigma})
\end{aligned}$$

with priors: $\alpha \sim \Gamma(\alpha; a_\alpha, b_\alpha(\text{rate}))$, $\boldsymbol{\mu} \sim N_2(\boldsymbol{\mu}; a_\mu, B_\mu)$, and $\boldsymbol{\Sigma} \sim IWish(\boldsymbol{\Sigma}; a_\Sigma, B_\Sigma)$, where $f(\boldsymbol{p}|\alpha) = \alpha^{L-1} p_L^{\alpha-1}(1-p_1)^{-1}(1-(p_1+p_2))^{-1} \times ... \times (1-\sum_{l=1}^{L-2} p_l)^{-1}$ is a special case of the generalized Dirichlet distribution as is Connor and Mosimann Connor & Mosimann (1969). Let n^* be the number of distinct components of $\mathbf{w}$ where $\mathbf{w}^* = \{\mathrm{w}_j^* : j = 1, ..., n^*\}$ are the distinct components. Let Ψ represent the vector of the most recent iteration of all other parameters. For $i = 1, ..., n$, let $\delta_i = 0$ if t_i is observed and $\delta_i = 1$ if t_i is right censored. Finally, let $b = 1, ..., B$ be the number of iterations in the MCMC. Then B samples from the joint posterior distribution, $p(\boldsymbol{\mu}, \boldsymbol{\sigma}^2, \mathbf{w}, \boldsymbol{p}, \lambda, \tau^2, \rho, \alpha)|data)$ are obtained for $b = 1, ..., B+1$:

Sample from the posterior conditional distribution for $\boldsymbol{\theta}_l$ for $l = 1, ..., L$: If l is not already a component: $l \notin \{\mathrm{w}_j^{*(b)} : j = 1, ..., n^{*(b)}\}$

$$p(\theta_l^{(b+1)}, \phi_l^{(b+1)}|data, \Psi) \overset{\text{draw}}{\sim} N_2(\boldsymbol{\mu}^{(b)}, \boldsymbol{\Sigma}^{(b)})$$

If l is an active component: $l \in \{\mathrm{w}_j^{*(b)} : j = 1, ..., n^{*(b)}\}$

$$p(\theta_l, \phi_l | data, \Psi) \propto N_2((\theta_l, \phi_l)'; \boldsymbol{\mu}, \boldsymbol{\Sigma}) \prod_{\{i:l=\mathrm{w}_i\}} \left[\Gamma(t_i; e^{\theta_l}, e^{\phi_l})\right]^{1-\delta_i} \left[\int_{t_i}^{\infty} \Gamma(u_i; e^{\theta_l}, e^{\phi_l}) du_i\right]^{\delta_i}$$

We use a Metropolis-Hastings step for this update. We sample from the proposal distribution $(\theta_l', \phi_l')' \sim N_2((\theta_l^{(b)}, \phi_l^{(b)})', cS^2)$, where S^2 is updated from the average posterior samples of $\boldsymbol{\Sigma}$ under initial runs, and $c > 1$. Draw $\eta \sim Unif(0, 1)$.

If $\eta < min \left\{ 1, \frac{N_2((\theta_l', \phi_l')'; \boldsymbol{\mu}^{(b)}, \boldsymbol{\Sigma}^{(b)}) \prod_{\{i:l=\mathrm{w}_i^{(b)}\}} \left[\Gamma(t_i; e^{\theta_l'}, e^{\phi_l'})\right]^{1-\delta_i} \left[\int_{t_i}^{\infty} \Gamma(u_i; e^{\theta_l'}, e^{\phi_l'}) du_i\right]^{\delta_i}}{N_2((\theta_l^{(b)}, \phi_l^{(b)})'; \boldsymbol{\mu}^{(b)}, \boldsymbol{\Sigma}^{(b)}) \prod_{\{i:l=\mathrm{w}_i^{(b)}\}} \left[\Gamma(t_i; e^{\theta_l^{(b)}}, e^{\phi_l^{(b)}})\right]^{1-\delta_i} \left[\int_{x_i}^{\infty} \Gamma(u_i; e^{\theta_l^{(b)}}, e^{\phi_l^{(b)}}) du_i\right]^{\delta_i}} \right\}$

set $(\theta_l^{(b+1)}, \phi_l^{(b+1)})' = (\theta_l', \phi_l')'$

else $(\theta_l^{(b+1)}, \phi_l^{(b+1)})' = (\theta_l^{(b)}, \phi_l^{(b)})'$.

Turning to the update for $\boldsymbol{p}$, we have:

$$p(\boldsymbol{p} | data, \Psi) \propto f(\boldsymbol{p}|\alpha) \prod_{l=1}^{L} p_l^{M_l} \; M_l = |\{i : \mathrm{w}_i = l\}| \; , l = 1, ..., L$$

$$p(\boldsymbol{p}^{(b+1)} | data, \Psi) \overset{\text{draw}}{\sim} \text{Generalized Dirichlet Distribution}$$

To sample from this distribution, for $l = 1, ..., L$ draw latent variable: $V_l^{*(b+1)} \overset{ind}{\sim} Beta(1 + M_l^{(b)}, \alpha^{(b)} + \sum_{r=l+1}^{L} M_r^{(b)})$. Now set $p_1^{(b+1)} = V_1^{*(b+1)}, p_l^{(b+1)} = V_l^{*(b+1)} \prod_{r=1}^{l-1} (1 - V_r^{*(b+1)})$ $(l = 2, ..., L-1)$, and $p_L^{(b+1)} = 1 - \sum_{l=1}^{L-1} p_l^{(b+1)}$.

Now we update w_i for $i = 1, .., n$:

$$p(\mathrm{w}_i^{(b+1)} | data, \Psi) \overset{\text{draw}}{\sim} \sum_{l=1}^{L} \tilde{p}_{li} \delta_{(l)}(\cdot)$$

where $\tilde{p}_{li} = \frac{p_l^{(b+1)} \left[\Gamma(t_i; e^{\theta_l^{(b+1)}}, e^{\phi_l^{(b+1)}})\right]^{1-\delta_i} \left[\int_{t_i}^{\infty} \Gamma(u_i; e^{\theta_l^{(b+1)}}, e^{\phi_l^{(b+1)}}) du_i\right]^{\delta_i}}{\sum_{l=1}^{L} p_l^{(b+1)} \left[\Gamma(t_i; e^{\theta_l^{(b+1)}}, e^{\phi_l^{(b+1)}})\right]^{1-\delta_i} \left[\int_{t_i}^{\infty} \Gamma(u_i; e^{\theta_l^{(b+1)}}, e^{\phi_l^{(b+1)}}) du_i\right]^{\delta_i}}, \; l = 1, ..., L.$

For the update for $\boldsymbol{\mu}$ we have:

$$p(\boldsymbol{\mu} | data, \Psi) \propto \prod_{\{l \in \mathbf{w}^*\}} N_2((\theta_l, \phi_l)'; \boldsymbol{\mu}, \boldsymbol{\Sigma}) N_2(\boldsymbol{\mu}; a_\mu, B_\mu)$$

$$p(\boldsymbol{\mu}^{(b+1)} | data, \Psi) \overset{\text{draw}}{\sim} N_2(m_\mu, S_\mu^2)$$

where $m_\mu = S_\mu^2 \left(B_\mu^{-1} a_\mu + \boldsymbol{\Sigma}^{-1} \sum_{\{l \in \mathbf{w}^{*(b+1)}\}} \boldsymbol{\theta}_l^{(b+1)}\right)$, $S_\mu^2 = \left(B_\mu^{-1} + n^{*(b+1)} \boldsymbol{\Sigma}^{-1(b)}\right)^{-1}$.

Turning to the update of $\boldsymbol{\Sigma}$, we have:

$$p(\boldsymbol{\Sigma} | data, \Psi) \propto \prod_{\{l \in \mathbf{w}^*\}} N_2((\theta_l, \phi_l)'; \boldsymbol{\mu}, \boldsymbol{\Sigma}) IWish(\boldsymbol{\Sigma}; a_\Sigma, B_\Sigma)$$

$$p(\boldsymbol{\Sigma}^{(b+1)} | data, \Psi) \overset{\text{draw}}{\sim} IWish(n^{*(b+1)} + a_\Sigma, B_\Sigma + \sum_{\{l \in \mathbf{w}^{*(b+1)}\}} (\boldsymbol{\theta}_l^{(b+1)} - \boldsymbol{\mu}^{(b+1)})(\boldsymbol{\theta}_l^{(b+1)} - \boldsymbol{\mu}^{(b+1)})')$$

Lastly, the update for α is given by:

$$p(\alpha | data, \Psi) \propto \Gamma(\alpha; a_\alpha, b_\alpha) f(\mathbf{p}|\alpha)$$

$$p(\alpha^{(b+1)} | data, \Psi) \overset{\text{draw}}{\sim} \Gamma\left(L + a_\alpha - 1, -\sum_{s=1}^{L-1} log(1 - V_s^{*(b+1)}) + b_\alpha\right)$$